

Evaluation of Multigrid Solutions for Turbulent Flows

Boris Diskin* and Hiroaki Nishikawa†

National Institute of Aerospace, Hampton, VA 23666

A multigrid methodology has been recently developed in a NASA solver, FUN3D, and successfully applied for a wide range of turbulent flows, from simple two-dimensional geometries to realistic three-dimensional configurations. The methodology is applicable to structured- and unstructured-grid solutions and includes both regular and agglomerated coarse meshes. Significant speed-ups over single-grid computations have been demonstrated. In the current work, we report on a detailed evaluation of the solver performance in computing benchmark turbulent flows. For those benchmark computations, multigrid solutions are compared with the corresponding single-grid solutions in terms of time-to-solution characteristics measured in the same computing environment. Multigrid efficiency enables a detailed grid-refinement study of a turbulent flow around an airfoil that is reported in the Appendix.

I. Introduction

The accuracy, efficiency, and robustness of turbulent flow computations are the critical factors to increased use of computational fluid dynamics (CFD) within the design process of aerospace vehicles. Current efforts are directed toward extending the scope of application of CFD methods to more complex and challenging turbulent and transitional flows with significant areas of separation. Among the active research directions are improved methods for modeling turbulence and transition, effective and automated methods for grid generation and adaptation, and highly accurate discretization schemes. While these research directions promise significant improvements in future practices, the current state of the art is represented by second-order finite-volume solvers for Reynolds-Averaged Navier-Stokes (RANS) equations with one or two-equation turbulence models. Recent Drag and High Lift Prediction Workshops^{1,2} demonstrated the advanced capability of the state-of-the-art solvers and the need for further improvements to adequately address increasingly complex turbulent flows.

This paper reports on recent advancements in developing and evaluating the “best-practice” methods in a state-of-the-art general-purpose unstructured CFD code, FUN3D,^{3,4} developed and supported at NASA Langley. The following four desirable properties guide the developments: (1) consistency of discrete approximations, (2) discrete solvability, (3) automation, and (4) efficiency. *Consistency of discrete approximations* implies that all discrete approximations to integro-differential operators possess some order property on all grids used in the solution process. One way to verify consistency (and the order property) of discrete approximations is to demonstrate convergence of discrete solutions to a continuous one on consistently refined grids.⁵ The importance of consistent grid refinement is often overlooked in current practices and may lead to misinterpretation of computational results. *Discrete solvability* requires the solver to be able to converge to ‘machine zero’ residuals for all equations (including the turbulence model equations) on all grids used in the solution process. While there is no consensus definition of ‘machine zero’ residuals, one possible definition is a norm of the residual response on a random solution perturbation with a relative amplitude of 10^{-15} . *Automation* implies that various algorithmic choices required for adaptive solution strategies are done automatically, without user intervention. *Efficiency* is defined as the solver ability to exploit massively parallel high-performance computational environments, advanced convergence acceleration methods and adaptation techniques to minimize time to solution.

*Research Fellow (bdiskin@nianet.org) National Institute of Aerospace (NIA), 100 Exploration Way, Hampton, VA 23666 USA.; Research Associate Professor, Mechanical and Aerospace Engineering Department, University of Virginia, Charlottesville, Associate Fellow AIAA

†Associate Research Fellow (hiro@nianet.org), National Institute of Aerospace 100 Exploration Way, Hampton, VA 23666 USA, Senior Member AIAA

In recent years, the FUN3D solver underwent several critical improvements advancing toward those desirable properties. Improved iteration methods employ better linearizations, efficient strategies for Jacobian computations and updates, adaptive pseudo-time stepping and under-relaxation, Krylov subspace methods, primal and agglomeration multigrid, and approximate-mapping methods. Improved discretization methods use more accurate gradient approximations, highly-stable and accurate diffusion schemes, and consistent discretizations on agglomerated grids. Improved scalability and high-performance computing characteristics are achieved by implementing asynchronous input/output methods, distributed scalable preprocessing, and advanced cache memory management, as well as by optimizing implementations of individual routines.

This paper is a sequel to a series of previous papers⁶⁻¹⁰ documenting development of multigrid techniques in FUN3D. An agglomeration multigrid (AgMG) method has been described in detail in Ref. [10]. The method used the Full Multigrid (FMG) scheme¹¹ to obtain initial solutions and nonlinear Full Approximation Scheme (FAS) multigrid cycle¹¹ to accelerate convergence. The computational results were presented for a series of 2D and 3D turbulent-flow benchmarks. The AgMG solver was compared with a state-of-the-art single-grid (SG) solver and with a structured multigrid (StMG) approach, which is referenced in this paper as a Primal Multigrid (PrMG) approach. A significant convergence acceleration was demonstrated in comparison to convergence of the SG solver. Convergence rates and number of iterations to convergence were similar between AgMG and PrMG schemes at CFL numbers bounded by 200. PrMG computations with larger CFL (up to 10,000) indicated a potential for an additional order of magnitude reduction in operation count and in time to solution.

In this paper, the following new developments are presented. The relaxation scheme is modified to allow AgMG and SG computations with large CFL numbers, leading to improved convergence for multigrid and single-grid iterations. The AgMG, PrMG, and SG solvers reported in this paper employ consistent discrete approximations (e.g., no edge-term-only approximations, sometimes referenced as thin-layer approximations, are used for mean-flow and turbulence-model equations) on all grids and satisfy discrete solvability and automation requirements. The main focus of this paper is on efficiency evaluation. The time to solution is now directly assessed for computations of two-dimensional (2D) and three-dimensional (3D) turbulent flows suggested at the Turbulence Modeling Resource (TMR) website¹² supported at NASA Langley Research Center and for turbulent flows around 3D configurations used in the Drag Prediction Workshop (DPW) series.^{1,13} For 2D test cases, the total time to obtain an accurate solution on a fine grid is reported. For 3D problems, the time to solution is evaluated using the work units computed with a CFD benchmark code TauBench¹⁴ developed at the German Aerospace Center (DLR).

The algorithmic and implementation issues are also discussed and illustrated by numerical examples. In particular, turbulent-flow solutions on simplex-element (i.e., 2D triangular and 3D tetrahedral) grids are revisited with new solver technologies such as the new version of the Spalart-Allmaras (SA) model as described at the TMR website,¹² including the provisions to accommodate negative turbulence¹⁵ and high-order convection approximations. Accuracy of turbulent flow solutions on simplex-element grids is significantly improved. Efficient multigrid solutions are reported for the 3D RANS formulations with a second-order discretization of the SA turbulence model.

Unusual performance of iterative solvers has been observed and documented for some test cases. For example, convergence of all iterative solvers slows down for 2D turbulent flows around NACA-0012 airfoil at zero angle of attack. Another example is a remarkable convergence of SG turbulent-flow solutions on structured grids around a hemisphere cylinder; the SG solver demonstrates time to solution comparable with multigrid solutions on the same grids.

The paper is organized as follows. First, methods used for evaluation of solver efficiency are described. Then, components of multigrid solvers are presented. Discussion of various implementation issues affecting performance of the current solvers is followed by the section presenting numerical tests for a 2D flat plate, a 2D NACA-0012 airfoil, a 3D Hemisphere Cylinder, and 3D wing-body DPW configurations. The final section contains concluding remarks. The appendix presents a grid-refinement study for a flow around a 2D NACA-0012 airfoil.

II. Evaluation of Solver Efficiency

In this paper, solver efficiency is assessed in terms of time to solution. For 2D computations, the time is measured as the wall clock time in seconds. For 3D computations, time is measured in work units; one work unit is the wall-clock time for running a benchmark code, TauBench,¹⁴ in the same computer environment as the evaluated code. Using the same benchmark code to compute a work unit should eliminate (or at least

reduce) the dependence on specific hardware platforms and compiler options.

TauBench is a light-weight C code that can be run on a cluster of parallel machines via MPI and provide a non-dimensional time (work unit) that is comparable on different machines and can be used as a basis for code run time comparisons. It was developed as a benchmarking tool that performs typical CFD operations and models the run time behavior of an explicit finite volume code. Among other things, the code simulates point and face loops and ghost point communication via MPI. It was successful in predicting performances of a large-scale CFD code, TAU, developed and supported at DLR on different parallel machines and showed a better correlation to TAU run times than other benchmarks.

In the tests reported in this paper, TauBench runs on the same processors with the same compiler and optimization flags as the evaluated solvers, approximately matching the number of grid points per processor and performing 10 time steps. Figure 1 shows the variation of the TauBench work unit in seconds versus the number of grid points per processor in computations on 192 processors. This plot characterizes the performance of the local computing cluster and predicts scalable computations on 192 processors with partitions between 5K and 500K of grid points per processor. Ten independent tests were performed for each partition size. Different processors were allocated for each test on a computational cluster shared by many users, so the system load varied in each test as well. The same processors were used for all partitions within a test. For a given partition size, the marker shows the mean TauBench computational time and the bar shows the limits of the time variation. Overall, the time is fairly consistent varying less than 16% for sparse partitions (small number of nodes per processor) and less than 4 % for dense partitions (many nodes per processor). Although not shown, two tests required noticeably larger times for computations on all partitions than other tests. This discrepancy indicates that similar processors do not perform identically and explains the location of the mean time marks at the lower parts of the variation bars. In the tests reported below, TauBench and the actual code always run on the same processors; the work unit time is taken as the TauBench execution time on those specific processors.

III. Multigrid Solvers

A. Primal Grids and Discretizations

A grid is defined as primal if it has the following elements: quadrilaterals and triangles in 2D; hexahedra, prisms, pyramids, and tetrahedra in 3D. The discretization method used here is a finite-volume discretization centered at nodes. It is based on the integral form of governing equations of interest. The governing equations are the RANS equations using an SA one-equation turbulence model.¹⁶ The primal-grid discretization of the viscous terms is obtained by the Green-Gauss (GG) scheme.^{3,17} For non-simplex cells, an edge-based augmentation is used to increase the h -ellipticity of the viscous operator.^{3,17} This augmentation is done by the face-tangent construction;⁸ thus the resulting scheme is called here the face-tangent GG scheme. The inviscid terms are discretized by a standard edge-based method with unweighted least-squares (LSQ) gradient reconstruction and Roe's approximate Riemann solver.^{18,19}

The turbulence model is the SA model as described at the TMR website,¹² including the provisions to accommodate negative turbulence.¹⁵ The convective term of the turbulence-model equation can be discretized with first or second-order accuracy while the diffusion terms are discretized with second-order accuracy by the face-tangent GG scheme. The node-centered velocity gradients used in the turbulence-model source evaluation are computed by the edge-based GG scheme on dual elements.

For high-aspect-ratio grids typical for high-Reynolds-number flows, grid nodes are grouped to form implicit lines typically extended from the body surface. This line information is used by discretization, implicit-relaxation, and agglomeration/coarsening procedures.

B. Primal Multigrid Solver

Primal grids used in the PrMG solver are regular structured grids, which are composed of blocks that can be mapped onto a Cartesian grid. A PrMG grid with non-Cartesian cells, i.e., cells other than 2D quadrilaterals and 3D hexahedra, can be derived from a Cartesian grid by regular division of corresponding 2D quadrilaterals into 2 triangles and 3D hexahedra into two prisms or six tetrahedra. Other types of primal grids, e.g., fully unstructured grids, have to use the AgMG solver for multigrid convergence acceleration. The coarse PrMG grids are nested and recursively derived from a finer grid by full coarsening, i.e., eliminating every other grid line/plane in each direction. The non-Cartesian cells on coarse grids are introduced analogously to corresponding cells on the finer grid. The primal-grid discretizations are used on all coarse grids. The

Table 1. Summary of discretizations used on primal and agglomerated grids.

Grids	Equations	Inviscid	Diffusion	Source
Primal	Mean flow	Second order	Face-tangent GG	
	Turbulence	First or second order	Face-tangent GG	Edge-based GG
Agglomerated	Mean flow	First order	Avg-LSQ	
	Turbulence	First order	Avg-LSQ	WLSQ

prolongation operator is a linear interpolation. On all 2D PrMG grids, the interpolation is bi-linear. On 3D hexahedral primal grids, the interpolation is tri-linear. On 3D prismatic primal grids used in Section IV.C the interpolation is either tri-linear or bi-linear within quadrilateral cell faces and edge-based linear within triangular cell faces. The residual restriction operator is fully weighted, globally conservative and constructed as a (scaled) transposition to the prolongation operator. Within a multigrid FAS cycle, the finer grid solution is restricted to the coarser grid by injection.

C. Agglomerated Grids and Discretizations

Hierarchical Agglomeration Scheme

As described in the previous papers,^{6–8,10} the grids are agglomerated within a topology-preserving framework, in which hierarchies are assigned based on connections to the computational boundaries: corners, ridges, valleys, and interiors. The agglomerations proceed hierarchically from seeds within the topologies — first corners, then ridges, then valleys, and finally interiors. Hierarchies on each agglomerated grid are inherited from the finer grid.

Full-Coarsening Line-Agglomeration for Viscous Grids

For viscous grids, the agglomeration is performed in the following sequence:

1. Agglomerate viscous boundaries (bottom of implicit lines, in general).
2. Agglomerate prismatic layers through the implicit lines (line agglomeration).
3. Agglomerate the rest of the boundaries.
4. Agglomerate the interior.

The second step is a line-agglomeration step where volumes are agglomerated along implicit lines starting from the volume directly above the boundary volume, thus preserving the boundary agglomerates.

Coarse-Grid Discretization

On coarse grids, the governing equations are directly discretized. For inviscid terms in coarse-grid discretizations, a first-order edge-based scheme is employed. The gradients at the agglomerated cells are computed using either LSQ or inverse-distance weighted least-squares scheme (WLSQ). The WLSQ gradient is used to compute the source term in the turbulence-model equation. For the mean-flow and turbulence diffusion terms, the Average-Least-Squares scheme (Avg-LSQ) with the face-tangent augmentation is employed.^{8,20,21} Table 1 shows a summary of discretizations.

Parallel Implementation

For parallel implementation, the hierarchical agglomeration algorithm is applied independently in each partition. Consequently, no agglomeration is performed across partition boundaries. No special modification is necessary for the line agglomeration as our partitioning guarantees that all nodes in each grid-line belong to the same partition (i.e., lines are partitioned in the form of condensed-nodes each of which contains all nodes in a line). Such an approach simplifies implementation, but generally limits the depth of the multigrid cycle. Implementation of the PrMG solver has fewer limitations on the depth of the multigrid cycle.

D. Prolongation and Restriction Operators in Agglomerated Multigrid Solver

The residuals are restricted to the coarser grid by summing the finer-grid residuals over the control volumes agglomerated into a coarser control volume. The solution is restricted to the coarser grid by taking the volume-weighted average. The correction to the finer grid is prolonged by the linear interpolation from a tetrahedron defined on the coarse grid.^{7,9} For volumes having negative coefficients in the linear interpolation, as typically happens near curved viscous boundaries, piecewise constant prolongation is performed.

E. Relaxation Scheme

The relaxation scheme is described in detail in Ref. [10]. Here only an abbreviated description is provided. The scheme is an hierarchical nonlinear iteration scheme. On the innermost level it uses a preconditioner based on a defect-correction method and iterates on a simplified first-order Jacobian. One preconditioner iteration involves an implicit-line pass through the portion of the domain where implicit grid lines are defined, followed by a point-implicit sweep through the entire domain. The number of preconditioner iterations may vary for different nonlinear iterations. This variable preconditioner is wrapped with a Generalized Conjugate Residual (GCR) method to form a Jacobian-free linear solver that uses Frechet derivatives to approximate linear residuals. A nonlinear controller assesses the correction computed by the linear solver. The controller is responsible for the CFL adaptation strategy and for deciding on necessity of the Jacobian update. As a result of this assessment, the suggested correction can be applied fully, partially, or completely discarded; the current Jacobian may be updated or reused in the next iteration; and the current CFL number may increase, decrease, or stay the same. In the current relaxation scheme, the iterations are loosely coupled, i.e., applied separately to the mean-flow and turbulence-model equations. Note that the multigrid iterations are coupled in the sense that the mean-flow and turbulence-model equations are solved on coarse grids and the mean-flow and turbulence variables use a coarse-grid correction. Initially, the CFL number is ramped over a prescribed number of iterations, but then it changes fully automatically within prescribed bounds. The coarse-grid corrections are also assessed by the nonlinear controller and can be applied fully, partially, or completely discarded.

IV. Implementation Issues

There is a large body of anecdotal and empirical evidence²² that accuracy of turbulent-flow solutions on fully simplex-element grids lags behind accuracy on mixed element grids with prismatic elements within boundary layers and comparable degrees of freedom. However, grids with non-simplex elements impose some penalties in the form of a larger operation count in computing residuals and forming Jacobians and, as a result, larger time per iteration. Simplex-element grids are also more flexible for anisotropic grid adaptation than mixed-element grids.⁴ It would be beneficial to find ways to improve solution accuracy on simplex-element grids.

Excellent efficiency of multigrid solvers reported in this paper allows comparison of accuracy of turbulent-flow solutions on very fine grids composed of elements of different types. Another variation parameter is the approximation order of the turbulence-model discretization, first and second-order discretizations for the turbulence convection term are considered. In a number of previous studies (see Ref. [23], for a 2D structured-grid example), it was noted that second-order discretizations of the turbulence model equations do not provide strong benefits for turbulent-flow solutions on 2D quadrilateral, 3D prismatic, and/or 3D hexahedral grids. This was in part the reason why first-order discretizations of the turbulence-model equation are used almost exclusively in the current state-of-the-art CFD codes. Another reason was that second-order discretizations of the turbulence-model equation may become non-positive and difficult to converge. It was deemed impractical to solve second-order discretizations of some traditional turbulence models.

Recent advancements in turbulence modeling¹⁵ and solver technology¹⁰ changed this perception and made possible convergence of RANS equations with second-order discretizations of the turbulence model. Preliminary results reported in this paper indicate that accuracy of turbulent flow solutions on unstructured simplex-element grids with second-order discretizations of the turbulence model significantly improved and became comparable with the accuracy of structured and mixed-element grid solutions. Nevertheless, there is an important limitation for use of multigrid solutions on simplex-element grids for turbulent flows. It was shown⁸ by analysis and numerical experiments that convergence of AgMG methods with line agglomeration deteriorates on high-aspect-ratio simplex-element grids.

In node-centered median-dual discretizations used in the PrMG, AgMG, and SG solvers reported in this

paper, the solution values for inviscid fluxes are reconstructed at the edge medians. The reconstruction scheme that uses solution and gradient at the node is referenced here as the $\kappa = 0$ scheme because, on Cartesian grids, it is analogous to the corresponding scheme from the κ -family.²⁴ It was found that averaging the $\kappa = 0$ reconstructed solution with the edge-averaged ($\kappa = 1$) solution improves robustness of iterative solvers in some cases. Such a reconstruction scheme is referenced as a $\kappa = 0.5$ scheme. Note that $\kappa = 0.5$ solutions are typically less accurate than corresponding $\kappa = 0$ solutions.

During the testing, some of implementation issues related to the built-in logic for CFL adaptation have been discovered. In the current implementation, the CFL number decreases if certain linear and/or non-linear residuals have not been sufficiently reduced in an individual iteration. When the residuals approach the machine-zero level, the target residual reduction cannot be achieved, and the CFL number unnecessarily decreases to a minimum value. In many cases, this decrease does not cause any problems. However, sometimes, this artificial CFL number decrease may lead to unexpected convergence patterns. For example, in the current implementation, the mean-flow and turbulence-model equations are weakly coupled in relaxation, and the corresponding CFL numbers are adapted independently. In some iterations, the residuals of the mean-flow equations can be several orders of magnitude lower than the residuals of the turbulence-model equation. In particular, the mean-flow residuals may approach the machine-zero level, and the mean-flow CFL number may decrease prematurely, while the turbulence-model residuals are still large, and the corresponding turbulence solution is still evolving. In this scenario, the mean-flow solution becomes practically frozen and cannot react on changes in the turbulence solution, which may lead to temporary increases of the mean-flow residuals well above the previously achieved nearly machine-zero level. In the future, the CFL adaptation logic will be improved to avoid unnecessary CFL decreases corresponding to solutions with residuals well below the requested tolerance.

Polar singularity has been previously identified as a solver issue for inviscid flux reconstruction;¹⁰ remedies such as an approximate mapping reconstruction^{25,26} and a WLSQ gradient reconstruction (the latter has some robustness limitations) have been suggested and successfully applied.¹⁰ In the current paper, polar singularity is treated with an alternative, simple, and robust method referenced as a line-mapping reconstruction. With this method, solution reconstruction in the directions aligned with the grid lines uses highly accurate one-dimensional gradients computed along those lines; reconstruction in all other directions is performed with the LSQ gradient.

Several other issues have been encountered during implementation of the “best-practice” methods. On some grids with implicit lines, the numbers of nodes in a line may differ significantly. For example, on structured NACA-0012 grids from the TMR website, the implicit lines in the wake region contain about twice as many nodes as the lines originated at the airfoil surface. For domain decomposition, a grid practitioner collapses implicit lines into super-nodes and assigns them weights, to ensure that partition boundaries do not interrupt implicit grid lines. In the current implementation, all the super-node weights are the same and may lead to a sub-optimal load balance. This implementation issue can be easily repaired.

It was discovered that use of string variables in conditional statements within a loop significantly slowed down the code performance; sometimes the execution time of the loop was more than doubled. Similar observations have been made elsewhere.²⁷ Conditional statements evaluating string variables were moved outside of the main loops.

In several cases, significant time variations (up to a factor of two) have been observed depending on specific computer nodes used in computations. All nodes used in these computations have identical specifications; compiler options in the runs are also identical. For a fair comparison, efforts have been made to use the same nodes in similar computations.

V. Numerical Results

A. 2D Flat Plate

This case is taken from the Turbulence Model Numerical Analysis section of the TMR website.¹² The geometry and flow as well as boundary conditions are illustrated in Figure 2. In comparison with the Flat Plate Verification case at the same website, the grids extend both upstream and downstream of the plate. This eliminates possible issues associated with prescribing pressure at the trailing edge of the plate.

A large sequence of nested quadrilateral grids of the same family are provided. The $\kappa = 0$ scheme is used for inviscid flux reconstruction. Implicit-lines are formed from the bottom to the top, covering the entire domain. Primal triangular grids are derived by regularly dividing primal-grid quadrilateral elements into

triangles. The AgMG solver is applied only to primal quadrilateral grids. A 4-level FAS(2, 2) cycle is used in both the PrMG and AgMG solvers. Initial approximations for multigrid solutions on the primal grids are obtained through the FMG algorithm with 10 cycles at each coarse grid. The initial approximation for the SG solver is obtained through a grid sequencing procedure that uses the PrMG coarse grids and performs 40 single-grid iterations on each coarse grid to approximately solve the coarse-grid equations. The CFL ramping strategy used on the coarsest grid in the FMG/grid-sequencing process increases the CFL number over the first ten cycles (40 single-grid iterations) from CFL = 10 to CFL = 100 for the mean-flow equations and from CFL = 20 to CFL = 200 for the turbulence-model equations. After these iterations, the CFL number is adapted automatically with the upper bound of CFL = 10,000. The initial CFL for the mean-flow and turbulence-model equations on all finer grids is the minimum of the corresponding CFL on the previous coarse grid and CFL = 200.

The convergence of the mean-flow and turbulence-model residuals on a 2561×769 quadrilateral grid is shown in Figure 3 to illustrate the number of cycles and time to solution. The L_2 norm of the mean-flow residuals (R:MF) is computed over the four mean-flow equations; the L_2 norm of the turbulence-model residuals (R:T) is computed over one turbulence-model equation. The case was run with 16 processors, leading to approximately 123K nodes per partition. The PrMG and AgMG solvers converged to the specified tolerance, 10^{-9} , with no major issues. The SG solver was aborted after approximately 3 hours of computations. At the moment of abort, the SG residuals were slowly converging; the mean-flow residuals were at the level of 10^{-11} and the turbulence-model residual was at the level of 10^{-8} . The multigrid solvers show an excellent efficiency and dramatically accelerate convergence in comparison with the SG solver. The multigrid solutions converge to the requested tolerance in a few cycles, reducing residuals by 5 orders of magnitude.

Figure 4 shows the history of CFL adaptation for the mean-flow and turbulence-model equations. For the first 42 FMG (12 fine-grid) cycles, the CFL number steadily increased according to the CFL adaptation strategy; the CFL numbers for all three solvers are overplotted. The CFL maximum of 10,000 was not reached in multigrid solutions because the convergence tolerance of 10^{-9} was met first. The turbulence-model CFL number in the SG solution reached the maximum in about 50 FMG cycles and remained at CFL = 10,000 for the rest of computations. After reaching the level of CFL = 5,000, the SG mean-flow CFL number oscillated between CFL = 200 and CFL = 5,000 and eventually stabilized at the level of CFL = 2,000. Figure 5 shows convergence of the drag and lift coefficients and confirms that multigrid solvers reduce the time for computing these aerodynamic quantities by an order of magnitude. For the AgMG solution, only finest-grid values of aerodynamic coefficients are shown.

Figure 6 shows convergence of the drag coefficient in grid refinement and illustrates the effects of mesh elements and accuracy of turbulence-model discretization. All solutions were computed with the PrMG solver. The first and second-order approximations of the turbulence-model equation differ in discretization of the convective term; the former uses a first-order discretization, while the latter uses a second-order discretization. Discretizations of the nonlinear diffusion and the source terms of the turbulence-model equation are the same in the first and second-order approximations. In agreement with earlier observations, second-order approximations of the turbulence-model equation have a little effect on the solution accuracy on quadrilateral grids,²³ and the solution accuracy on triangular grids with the first-order approximation is significantly worse than the solution accuracy on quadrilateral grids.²² However, a new observation is that the solution accuracy on triangular grids with the second-order approximation is dramatically improved and comparable to the accuracy on quadrilateral grids. Table 2 shows the time to convergence of the AgMG and PrMG solvers on quadrilateral and triangular grids with the first- and second-order discretizations of the convection term in the turbulence-model equations. Both solvers show uniformly excellent efficiency for all test cases.

B. NACA 0012 airfoil

Turbulent flows over the NACA 0012 airfoil at two angles of attack, $\alpha = 0^\circ$ and $\alpha = 10^\circ$, are considered. The cases are taken from the Turbulence Model Validation Cases and Grids section of the TMR website.¹² The computations are performed on the finest available C-type 1793×513 grid with the near-incompressible free-stream conditions recommended by the website (the Mach number is 0.15 and $Re = 6.0 \times 10^6 / (\text{unit grid length})$). The $\kappa = 0$ scheme is used for inviscid flux reconstruction. The far-field boundary is located about 500 chords away from the airfoil. The boundary conditions are the adiabatic solid wall on the airfoil surface and inviscid Riemann boundary conditions at the far-field boundary. The cases run on 16 processors with approximately 57K nodes per partition. A 4-level FAS(2, 2) cycle is used in the PrMG solver and a 3-level FAS(3, 3) is used in the AgMG solver. The depth of multigrid cycles is limited by the

Table 2. Summary of the time to convergence (in seconds) of multigrid solvers for 2D test cases is shown. In parenthesis is the number of finest-grid relaxations to convergence. Q-1 and Q-2 denote fully quadrilateral grids with the first- and second-order approximations for the convection term in the SA turbulence-model equation, respectively. T-1 and T-2 denote fully triangular grids with the first- and second-order approximations for the convection term in the SA turbulence-model equation, respectively. The AgMG and PrMG solvers run on 16 processors and converge to the residual tolerance of 10^{-9} . The CFL3D solver runs 10,000 W(1,1) cycles for NACA-0012 flows on a single processor for about 32 hours; for comparison, the time in the table is the total run time divided by 16. For the NACA-0012 $\alpha = 0^\circ$ flow, CFL3D residuals did not converge after 10,000 cycles. For the NACA-0012 $\alpha = 10^\circ$ flow, the final CFL3D density residuals were at the level of 10^{-13} , and the corresponding turbulence-model residuals were at the level of 10^{-8} .

Case		Flat Plate	NACA-0012 0°	NACA-0012 10°
Grid		2561×769	1793×513	1793×513
CFL3D	Q-1	–	not converged	7200(20,000)
AgMG	Q-1	365(40)	2102(882)	897(354)
	Q-2	450(40)	2359(762)	1198(378)
PrMG	Q-1	344(48)	1348(524)	434(176)
	Q-2	331(40)	1366(516)	470(196)
	T-1	307(32)	1724(656)	988(380)
	T-2	451(42)	1731(648)	converged to 10^{-8}

ability of the partitioner to balance the load without grid-line interruption.

Multigrid solvers generate initial solutions on primal quadrilateral and triangular grids through the FMG algorithm with 10 cycles at each coarser grid; the initial approximation for the SG solver is obtained through a grid sequencing procedure with 40 SG iterations on coarser grids. The CFL ramping strategy increases the CFL number over the first 50 FMG cycles (200 SG iterations) from $\text{CFL} = 20$ to $\text{CFL} = 200$ for both the mean-flow and turbulence-model equations. The CFL variation is effectively prescribed over all coarse-grid iterations and over several initial finest-grid iterations (10 PrMG cycles, 20 AgMG cycles, 40 SG iterations). After these iterations, the CFL number is adapted automatically with the upper bound of $\text{CFL} = 10,000$.

The convergence of the mean-flow and turbulence-model residuals is shown in Figures 7 and 8 for quadrilateral grids and in Figure 9 for triangular grids. The L_2 norm of the mean-flow residuals (R:MF) is computed over the four mean-flow equations; the L_2 norm of the turbulence-model residuals (R:T) is computed over one turbulence-model equation. As noted previously¹⁰ and confirmed by the current computations, the SG solver converges extremely slowly for the flow at $\alpha = 0^\circ$ angle of attack. Multigrid solvers significantly accelerate convergence, but still the multigrid performance for the flow at $\alpha = 0^\circ$ angle of attack is inferior in comparison with the performance for flows at larger angles of attack. While there is a lot of anecdotal confirmation of this behavior, the authors are not aware of any systematic study or a satisfactory explanation of this phenomenon. Note also that the increase of the mean-flow residuals observed at the final stage of the AgMG computations shown in Figures 7 (a) and 8 (a) is the result of a premature CFL number decrease discussed in Section IV. Figure 10 shows the history of CFL evolution for the mean-flow and turbulence-model equations. At the end of these computations, the mean-flow residuals are well below the requested tolerance of 10^{-9} .

Table 2 presents some quantitative characteristics of multigrid solutions. The PrMG and AgMG solvers converge to the specified tolerance, 10^{-9} , for almost all test cases. For the case of $\alpha = 10^\circ$, triangular grid, and second-order approximation to the convection term in the turbulence-model equation, the tolerance of 10^{-9} appears to be very close to the “machine-zero” residual. The PrMG solver quickly converges the mean-flow residuals to the level well below 10^{-9} and the turbulence-model residuals to the level of 10^{-8} and then the turbulence-model residual convergence becomes flat. Figures 9 (c) and (d) demonstrate that the SG solver behaves similarly after residuals reached similar levels of convergence.

For the flows computed on quadrilateral grids, the PrMG and AgMG solvers reduce the number of the fine-grid iterations and the total time to convergence by a factor of 7 or more in comparison with the SG solutions. As an additional reference, another CFD code, CFL3D,²⁸ was applied to the same NACA-0012 turbulent flows on the same quadrilateral 1793×513 grid. CFL3D is a well established second-order accurate structured multi-zone body-fitted cell-centered solver developed at NASA Langley and widely recognized for its efficiency. CFL3D solves the full Navier-Stokes formulation for the mean-flow equations and uses a

traditional strictly positive turbulence-model equation with a first-order approximation for the convection term and a thin-layer approximation for the diffusion term. The CFL3D iteration is a 3-level W(1, 1) multigrid cycle. On the considered 1793×513 grid, CFL3D ran 10,000 iterations on a single processor for about 32 hours starting from a converged solution for a nearby angle of attack. Assuming the perfect scalability of the solver, the optimum time on 16 processors would be about 2 hours or 7,200 seconds. After 10,000 iterations for the flow at $\alpha = 10^\circ$ angle of attack, the density residual norm was at the level of 10^{-13} , and the turbulence-model residual norm was at the level of 10^{-8} . As mentioned earlier, the case of the flow at $\alpha = 0^\circ$ angle of attack is more difficult to compute than other NACA-00012 cases. For this case, the CFL3D residuals did not converge after 10,000 iterations, but the drag coefficient appeared converged reasonably well. It is a widely accepted practice to terminate iterative convergence when the aerodynamic coefficients of interest appear sufficiently converged. Note, however, that there are many examples where a computed aerodynamic coefficient stays flat for many iterations, but eventually varies and converges to a significantly different final value. One of recent examples is DPW-5 computations reported in Ref. [10]. To avoid such an ambiguity, it is important to be able to converge discrete residuals to low levels approaching machine-zero residuals.

Figures 11 and 12 show convergence of the drag and lift coefficients. The AgMG solution is shown on the finest grid only. The PrMG solver on quadrilateral grids is about two times faster than the AgMG solver on these grids. Second-order discretizations of the convective term in the turbulence model equation have little effects on convergence.

Comparison of the discretization accuracy on quadrilateral and triangular grids for the NACA-0012 airfoil proved to be nontrivial. One of the main difficulties is to find a reliable estimate for a grid converged quantity of interest. A detailed grid-refinement study is reported in Appendix A for the case of $\alpha = 10^\circ$. The study shows that the drag coefficient can be extrapolated from relatively coarse grids to predict the grid-refined value. Figures 13 (a) and (b) show convergence of the drag coefficient on a family of quadrilateral and triangular grids for flows at $\alpha = 0^\circ$ and $\alpha = 10^\circ$, respectively. Quadrilateral grids are provided at the TMR website; triangular grids are derived from the corresponding quadrilateral grids by dividing each quadrilateral element into two triangles. First- and second-order approximations are used for the convection term in the turbulence-model equation. As expected, the drag coefficients computed on quadrilateral grids show little sensitivity to variations in this approximation order. The figures show that the drag coefficient computed on finer triangular grids with the second-order approximation for the turbulence convection term is close to the coefficients computed on quadrilateral grids. The drag coefficient computed on triangular grids with the first-order turbulence convection is significantly smaller. Note however, that similarity of drag coefficients computed on triangular and quadrilateral grids does not necessarily imply that all other functionals of interest are accurate (or similar) on those grids. Results presented in Appendix A show that accuracy and convergence of the lift and the pitching moment coefficients are much more sensitive to the local grid density near the trailing edge than accuracy and convergence of the drag coefficient. The lift and the pitching moment can change significantly between grids with similar elements and degrees of freedom, but with different spatial distributions of those degrees of freedom. Extremely dense grids are required to resolve important flow features and establish reliable estimates for grid converged quantities. While efficient multigrid solvers enable turbulent-flow computations on extremely dense 2D grids, grid adaptation has to be employed to improve accuracy of lift and moment computations for realistic large-scale 3D configurations.

C. 3D Hemisphere Cylinder

A 3D turbulent flow over a hemisphere cylinder is considered. This case is a new addition to the Turbulence Model Numerical Analysis section at the TMR website.¹² The geometry is taken from the experimental model studied in Ref.[29]. Figure 14 shows a schematic layout of the grids, along with boundary conditions. The free-stream flow parameters are Mach number $M = 0.6$, and the angle of attack $\alpha = 5.0^\circ$. In the experiment, $Re = 4.2 \times 10^6/ft$, the radius of the hemisphere was 0.5 inches, and the body length was 10 inches. Thus, $Re = 3.5 \times 10^5/(\text{unit grid length})$, with a hemisphere radius $R = 0.5$ and the body length $L = 10.0$ in the computational grid. The precise position and shape of the outer boundaries depend on the grid generation parameters. In the reported computations, the average distance from the surface is 80 radii.

Two families of primal grids are generated. Primal grids of Family I are based on hexahedral elements and use polar grids over the hemisphere. Near the polar axis, the hexahedral elements degenerate into triangular prisms, leading to polar singularity. The number of such prisms increases on finer grids. Line mapping is applied in the mean-flow inviscid reconstruction, improves accuracy and stability of the discrete equations, and enables convergence on grids with polar singularity. A primal hexahedral grid is illustrated in Figure 15.

The outer boundary is a hemisphere with the radius of 40 unit grid lengths.

Grids in Family II are generated to avoid polar singularity.¹⁰ Primal grids from this family are based on triangular prisms. The surface grid over the hemisphere is initially divided into six triangular sectors. Finer primal surface grids are generated through subdividing each triangular surface element into 4 smaller triangles by edge bisecting. This procedure guarantees that each interior surface node on an arbitrarily fine primal grid is edge-connected to exactly six other surface nodes. The shape of the outer boundary is chosen to facilitate generation of relatively isotropic grids in the far-field region of the computational domain and is illustrated in Figure 16. The fine primal grids of Family II are shown in Figures 17. Corresponding primal and agglomerated coarse grids are shown in Figures 18. No line mapping is applied.

For both families, implicit lines are formed throughout the entire domain, and the $\kappa = 0$ scheme is used for inviscid fluxes. The PrMG solver uses 4-level FAS(2, 1) and FAS(2, 2) cycles for primal grids of Families I and II, respectively. A 4-level FAS(3, 3) is used by the AgMG solver on prismatic grids of Family II. The SG solver is implemented as a 2-level FAS(1, 0) cycle, which does not apply any coarse-grid correction. An FMG method with 10 coarse-grid cycles is used to form initial approximations. The CFL number is ramped from CFL = 10 to CFL = 200 over 50 FMG cycles and the CFL bound is set at 10,000. The computations on primal grids with approximately 20M nodes are performed on 192 processors. In Figures 19 and 20, convergence of residuals and forces is shown for the AgMG, PrMG, and SG solvers on grids of the Families I and II, respectively. The L_2 norm of the mean-flow residuals (R:MF) is computed over the five mean-flow equations; the L_2 norm of the turbulence-model residuals (R:T) is computed over one turbulence-model equation. The residual tolerance is set to be 10^{-10} . The work units are computed with TauBench as explained in Section II. The TauBench execution time corresponding to 100K nodes per processor is 3.674 seconds. A first-order approximation to the turbulence convection term is used on grids of Family I; first and second-order approximations are used on grids of Family II.

On grids of Family I, solvers with second-order approximations to the turbulence convection did not converge to the required tolerance because the line mapping was not implemented for the turbulence-model equation. Figure 21 shows evolution of residuals and forces in PrMG computations with second-order turbulence convection on grids of Family I. Although convergence of the mean-flow and turbulence-model residuals stalls on all grids, the values of the aerodynamic coefficients appear to be very similar to the values computed from the converged solutions on grids of Families I and II (see Figures 19 and 20). Figure 22 shows that the largest residuals and the negative turbulence working variables are located on the polar axis on the finest 20M grid of Family I. It is expected that implementing the line-mapping method for inviscid reconstruction in the turbulence-model equation will enable convergence on grids of Family I for formulations with second-order turbulence convection.

Table 3 presents a summary of multigrid performance on grids with different elements. The PrMG solver demonstrates an excellent efficiency on grids of Family I with first-order turbulence convection. About four orders of magnitude in residual reduction is achieved in 17 FAS(2, 1) cycles (51 relaxations). The total time to solution on the 20M grid is 214 work unit or about 788 seconds. The SG solver on the same grid is much slower; it requires 144 iterations to converge that takes 1,186 work units or 4,356 seconds.

The PrMG solver is almost three times faster than the other two solvers on grids of Family II. One observation is that second-order approximations to the turbulence convection term do not have adverse effects on the solver efficiency. Another observation is that the AgMG and SG solvers show similar efficiency, converging to the required tolerance in about 800 work units for both first- and second-order approximations to the turbulence convection.

The latter observation is somewhat surprising because similar AgMG and SG computations with a first-order approximation to turbulence convection have been reported previously¹⁰ on mixed-element grids with 1M, 4M, 7M, 10M, and 15M nodes and with the CFL bound of 200. The AgMG solver demonstrated grid-independent convergence, converged residuals by 10 orders of magnitude in about 80 FAS(3, 3) cycles and provided a significant speed-up over the SG solutions that slowed down (as expected) on finer grids.

In contrast to those fully unstructured computations, the performance of the AgMG solver on the current 20M grids of Family II is comparable to the performance of the SG solver. The reason is an unexpectedly fast convergence of the SG iterations, not a deteriorated performance of the AgMG solver. The time to convergence for the current AgMG solver is actually better than the time to convergence shown in Ref. [10] in spite of the fact that the current partitions are significantly coarser (104K vs 60K nodes per partition) and higher-order approximations to the turbulence convection terms are used. The SG solver with the CFL bound of 10,000 dramatically improves its performance over the corresponding solver with the CFL bound of 200. Although not shown, in the current computations, the CFL number steadily grows for all solvers,

Table 3. A summary of the time to convergence (in work units) of multigrid solvers for a turbulent flow on 3D hemisphere cylinder $20M$ grids with different elements and the first- and second-order approximations for the convective term in the turbulence-model equation. The AgMG solver uses a 4-level FAS(3, 3) cycle. The PrMG solver uses 4-level FAS(2, 1) and FAS(2, 2) cycles on grids of Families I and II, respectively. The SG solver is implemented as a 2-level FAS(1, 0) cycle, which does not apply any coarse-grid correction. All computations use 192 processors and converge to the turbulence-model residual tolerance of 10^{-10} . The work unit is taken as 3.674 seconds and represents a typical wall-clock time for a TauBench run with 100K nodes per processor. In parenthesis is the number of finest-grid relaxations to solution.

Solver	Order	Family I	Family II
AgMG	1	–	785(246)
	2	–	819(246)
PrMG	1	214(51)	380(112)
	2	–	417(124)
SG	1	1,186(144)	768(137)
	2	–	817(158)

indicating that the solvers do not experience convergence difficulties.

In particular, Table 3 shows that the numbers of nonlinear iterations used by the PrMG and SG solvers on the finest prismatic grids of Family II are close. For example, to converge to the residual tolerance of 10^{-10} with a first-order approximation to the turbulence convection on the $20M$ grid of Family II, the PrMG solver needs 28 FAS(2, 2) cycles, which is equivalent to 112 relaxations, and the SG solver needs 137 iterations. The difference in the convergence time between the PrMG and SG solvers is explained by a significantly greater number of subiterations required by the SG solver to achieve a prescribed reduction (typically 0.5) of the linear residual. The subiterations are terminated when this reduction has been achieved or when the number of subiterations has exceeded a prescribed maximum. The number of subiterations differs from iteration to iteration. The PrMG relaxation typically requires 10 subiterations to meet the residual reduction target; the SG solver often uses many dozens of subiterations, but nevertheless also always meets the reduction target. The AgMG solver requires about two times more nonlinear iterations than the PrMG and SG solvers, but uses about the same number of subiterations per relaxation as the PrMG solver. Convergence of the aerodynamic coefficients is illustrated in Figures 20 (c) and (d). On these plots, the AgMG and PrMG solvers require significantly less time than the SG solvers to converge lift and drag coefficients. The AgMG computations are shown on the finest grid only.

D. DPW Configurations

A turbulent flow over a wing-body-tail configuration³⁰ is considered with Mach number $M = 0.85$ and $Re = 5.0 \times 10^6 / (\text{unit grid length})$. This configuration is used as a benchmark in the Drag Prediction Workshop (DPW) series. Several fully unstructured mixed-element grids were generated for DPW-4.¹³ Results here are shown on the medium grid of approximately $10M$ nodes for the angle of attack of $\alpha = 2.5^\circ$. The L_2 norm of the mean-flow residuals (R:MF) is computed over five mean-flow equations; the L_2 norm of the turbulence-model residuals (R:T) is computed over one turbulence-model equation. The residual tolerance for R:MF and R:T is set to be 10^{-8} . The AgMG solver has 3 levels; primal and agglomerated grids are shown in Figure 23. FAS(3, 3) cycles and $\kappa = 0.5$ discretization for inviscid fluxes are used. The FMG method with 10 cycles on coarse grids is used to form the initial solution. The CFL ramping strategy is to increase CFL from 10 to 200 over the first 50 cycles, which include coarse grid cycles within the FMG process. The maximum CFL number is set to 10,000. The computations are done on 180 processors. The approximate partition size is 57K per processor, and the corresponding TauBench work unit is 2.473 seconds. The required tolerance is achieved in about 8,000 seconds and 25,000 seconds for multigrid solutions with the first and second-order approximations to turbulence convection terms, respectively.

This case is very challenging. In the previous effort,¹⁰ the authors were able to converge both mean-flow and turbulence-model residuals to the required tolerance only using an (inconsistent) edge-terms-only discretization for the turbulence diffusion term. In the current computations, all discretization components are consistent. The convergence is achieved due to improvements in the relaxation scheme that provide accurate Jacobian approximations through Frechet derivatives, effective nonlinear control of the solution

updates, and management of the adaptive CFL and Jacobian updates. The convergence of the mean-flow and turbulence-model residuals is shown in Figure 24 for AgMG and SG solvers of the formulations with first- and second-order approximations to the turbulence convection.

All solutions converge to the required tolerance. The single-grid solution with a second-order approximation converges very slowly and is only shown for 20,000 work units. The solutions with the second-order approximation take about three times longer compared with the corresponding solutions with the first-order approximations. The CFL history shown in Figure 25 indicates that both the AgMG and SG solvers experience numerous problems, but are able to recover and converge residuals to the required tolerance. The AgMG solver requires especially low $O(1)$ CFL for the mean-flow equations in some iterations in computations with the second-order approximation. On the other hand, the SG solver with the second-order approximation uses the maximum turbulence-equation $CFL=10,000$ in some iterations. The AgMG solution is about four times faster than the SG solution both in terms of residual convergence and in terms of convergence of aerodynamic coefficients. See Figures 26 (a) and (b) for convergence of drag and lift coefficients, respectively. Note that convergence of aerodynamic coefficients is similar with first- and second-order approximations to the turbulence convection.

Families of structured primal grids for a transonic NASA Common Research Model Wing-Body configuration³¹ are available at the DPW-5 website¹ and preprocessed into hybrid grids; prismatic grids extend from the surface outward a fixed percentage of the points in the normal direction and tetrahedra fill the remainder of the domain. A grid and a solution are illustrated in Figure 27. Flow conditions are Mach number $M = 0.85$ and $Re = 5.0 \times 10^6 / (\text{unit grid length})$. Inviscid fluxes are discretized with the $\kappa = 0$ discretization. The convective term in the turbulence-model equation is discretized with a first-order scheme. The PrMG solver uses a 4-level FAS(2, 2) cycle. The FMG solver solves the coarse-grid equation to the required tolerance. The CFL number is ramped from $CFL = 10$ to $CFL = 200$ over the first 50 FMG cycles; the upper CFL bound is 10,000. Convergence of the PrMG solver for the flow at $\alpha = 1^\circ$ is reported for the third coarsest workshop grid, which has approximately 5.2M grid points.

The multigrid convergence of the mean-flow and turbulence-model residuals is shown in Figure 28 (a). For this computation, the residual tolerance is defined for the collective L_2 norm computed over six equations. It is set to be either 10^{-7} in the absolute norm value or a reduction of 8 orders of magnitude from the maximum norm exhibited on the current grid. The PrMG solver efficiently converges on all grids reducing residuals by eight orders of magnitude on all grids in the FMG process. The computation is performed on 192 processors, which roughly corresponds to 27K nodes per processor, and takes about 2,473 seconds and 207 (828) fine-grid cycles (relaxations) to meet the tolerance. The corresponding TauBench work unit is about 1 second. The CFL history is shown in Figure 28 (b). On all grids, the mean-flow residuals converge to much lower levels than the turbulence-model residuals, which is typical. However, the mean-flow CFL number is reduced more frequently, which is less typical. Convergence of the lift and drag coefficients is shown in Figure 29. The residuals of the turbulence-model equation are also shown to distinguish between forces computed on different grids. Overall, the lift and drag coefficients converge quickly to their final values on each grid and demonstrate apparent second-order convergence in grid refinement.

Convergence of residuals and forces in the PrMG and SG computations is compared in Figure 30. The SG residuals did not converge to the required tolerance after more than 22,858 seconds and 8,000 fine-grid iterations. The aerodynamic forces converge to the same values in both the PrMG and SG solutions, but the multigrid solution approaches the final values for the drag and lift coefficients more than five times faster.

VI. Conclusions

The research presented in this paper is work in progress. The multigrid methodology for turbulent-flow solutions developed and reported in previous papers has been enhanced and evaluated in terms of time to solution on a series of test cases ranging from simple two-dimensional (2D) flows to some realistic three-dimensional (3D) configurations. The computations have been shown with new solver technologies such as the version of the Spalart-Allmaras (SA) model including the provisions to accommodate negative turbulence and a high-order convection approximation in the turbulence-model equation. It was shown that the primal multigrid (PrMG) and agglomeration multigrid (AgMG) solvers significantly reduced time to solution in comparison with the single-grid (SG) iterations for all reported cases, including realistic 3D turbulent flows on fully unstructured grids with a second-order discretization of the convection term in the SA turbulence model. Second-order approximations of the turbulence convection improved accuracy of turbulent-flow solutions on simplex-element grids. The PrMG solver was the fastest solver in all reported

computations on structured grids. The AgMG solver showed comparable performance on structured grids and provided a significant speed-up over the SG solver on unstructured grids.

The multigrid efficiency enables a detailed consistent grid refinement study for a turbulent flow around the NACA-0012 airfoil on grids containing up to 16 million nodes. This study has revealed that accuracy of the extrapolated estimates for the grid-refined values of the lift and the pitching moment coefficients is extremely sensitive to local grid density; accuracy of the projected value for the drag coefficient is less sensitive to grid density. The apparent convergence of the forces and moments exhibited on coarse grids may lead to erroneous conclusions because the convergence trends may change dramatically on finer grids. Improved discretizations and grid adaptation methods will be required to perform similar studies for 3D turbulent flows in complex geometries.

Acknowledgments

The authors acknowledge support for this research from the National Institute of Aerospace. The first author acknowledges the support from the NASA cooperative agreement #NNL09AA00A. All results presented were computed with the FUN3D and CFL3D codes at NASA Langley Research Center. The authors are deeply grateful for help provided by Drs. James L. Thomas and Christopher L. Rumsey, NASA Langley Research Center. Dr. Thomas encouraged and supported this study in many ways and Dr. Rumsey helped with grid generation, CFL3D computations, and all aspects related to the TMR website.

References

- ¹“5th AIAA CFD Drag Prediction Workshop Website,” <http://aac.larc.nasa.gov/tsab/cfdlarc/aiaa-dpw>, January, 2014.
- ²“2nd AIAA CFD High Lift Prediction Workshop (HiLiftPW-2) Website,” <http://hiliftpw.larc.nasa.gov>, January, 2014.
- ³Anderson, W. K. and Bonhaus, D. L., “An implicit upwind algorithm for computing turbulent flows on unstructured grids,” *Comput. Fluids*, Vol. 23, 1994, pp. 1–21.
- ⁴Park, M. A., Laffin, K. R., Chaffin, M. S., Powell, N., and Levy, D. W., “CFL3D, FUN3D, and NSU3D Contributions to the Fifth Drag Prediction Workshop,” *51st Aerospace Sciences Meeting*, AIAA Paper 2013-0050, Grapevine, Texas, 2013.
- ⁵Thomas, J. L., Diskin, B., and Rumsey, C. L., “Towards Verification of Unstructured Grid Methods,” *AIAA J.*, Vol. 46, No. 12, December 2008, pp. 3070–3079.
- ⁶Nishikawa, H., Diskin, B., and Thomas, J. L., “Critical Study of Agglomerated Multigrid Methods for Diffusion,” *AIAA J.*, Vol. 48, No. 4, April 2010, pp. 839–847.
- ⁷Nishikawa, H., Diskin, B., and Thomas, J. L., “Development and Application of Agglomerated Multigrid Methods for Complex Geometries,” *40th AIAA Fluid Dynamics Conference and Exhibit*, AIAA Paper 2010-4731, Chicago, 2010.
- ⁸Thomas, J. L., Diskin, B., and Nishikawa, H., “A Critical Study of Agglomerated Multigrid Methods for Diffusion on Highly-Stretched Grids,” *Comput. Fluids*, Vol. 41, No. 1, February 2011, pp. 82–93.
- ⁹Nishikawa, H. and Diskin, B., “Development and Application of Parallel Agglomerated Multigrid Methods for Complex Geometries,” *20th AIAA Computational Fluid Dynamics Conference*, AIAA Paper 2011-3232, Hawaii, 2011.
- ¹⁰Nishikawa, H., Diskin, B., Thomas, J. L., and Hammond, D. P., “Recent Advances in Agglomerated Multigrid,” *51st AIAA Aerospace Sciences Meeting*, AIAA Paper 2013-863, Grapevine, TX, 2013.
- ¹¹Brandt, A. and Livne, O., *Multigrid Techniques: 1984 Guide with Applications to Fluid Dynamics, Revised Edition*, SIAM, 2011.
- ¹²“Turbulence Modeling Resource Website,” <http://turbmodels.larc.nasa.gov>. January, 2014.
- ¹³“4th AIAA CFD Drag Prediction Workshop Website,” <http://aac.larc.nasa.gov/tsab/cfdlarc/aiaa-dpw/Workshop4>, January, 2014.
- ¹⁴“The TauBench Benchmark,” <http://www.ipacs-benchmark.org/index.php?s=download&unterseite=taubench>. January, 2014.
- ¹⁵Allmaras, S. R., Johnson, F. T., and Spalart, P. R., “Modifications and Clarifications for the Implementation of the Spalart-Allmaras Turbulence Model,” *Seventh International Conference on Computational Fluid Dynamics*, Big Island, Hawaii, 2012.
- ¹⁶Spalart, P. R. and Allmaras, S. R., “A One-Equation Turbulence Model for Aerodynamic Flows,” *Recherche Aerospatiale*, No. 1, 1994, pp. 5–21.
- ¹⁷Diskin, B., Thomas, J. L., Nielsen, E. J., Nishikawa, H., and White, J. A., “Comparison of Node-Centered and Cell-Centered Unstructured Finite-Volume Discretizations: Viscous Fluxes,” *AIAA J.*, Vol. 48, No. 7, July 2010, pp. 1326–1338.
- ¹⁸Roe, P. L., “Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes,” *J. Comput. Phys.*, Vol. 43, 1981, pp. 357–372.
- ¹⁹Diskin, B. and Thomas, J. L., “Accuracy Analysis for Mixed-Element Finite-Volume Discretization Schemes,” *NIA Report No. 2007-08*, 2007.
- ²⁰Nishikawa, H., “Beyond Interface Gradient: A General Principle for Constructing Diffusion Schemes,” *40th AIAA Fluid Dynamics Conference and Exhibit*, AIAA Paper 2010-5093, Chicago, 2010.

²¹Nishikawa, H., “Robust and Accurate Viscous Discretization via Upwind Scheme - I: Basic Principle,” *Comput. Fluids*, Vol. 49, No. 1, October 2011, pp. 62–86.

²²Rumsey, C. L., Slotnick, J. P., Long, M., Stuever, R. A., and Wayman, T. R., “Summary of the First AIAA CFD High-Lift Prediction Workshop,” *J. Aircraft*, Vol. 48, No. 6, 2011, pp. 2068–2079.

²³Morrison, J. H., “Numerical Study of Turbulence Model Predictions for the MD 30P/30N and NHLP-2D Three-Element Highlift Configurations,” Nasa/cr-1998-208967, 1998.

²⁴van Leer, B., “Upwind-Difference Methods for Aerodynamic Problems Governed by the Euler Equations,” *Fluid Mechanics, Lectures in Applied Mathematics*, Vol. 22, 1985, pp. 327–335.

²⁵Diskin, B. and Thomas, J. L., “Comparison of Node-Centered and Cell-Centered Unstructured Finite-Volume Discretizations: Inviscid Fluxes,” *AIAA J.*, Vol. 49, No. 4, 2011, pp. 836–854.

²⁶Diskin, B. and Thomas, J. L., “Comparison of Node-Centered and Cell-Centered Unstructured Finite-Volume Discretizations: Inviscid Fluxes. (Composition erratum),” *AIAA J.*, Vol. 51, No. 1, 2013, pp. 277–277.

²⁷Derlaga, J. M., Phillips, T. S., and Roy, C. J., “SENSEI Computational Fluid Dynamics Code : A Case Study in Modern Fortran Software Development,” AIAA Paper 2013-2450, 2013.

²⁸Krist, S. L., Biedron, R. T., and Rumsey, C. L., “CFL3D User’s Manual (Version 5.0),” NASA TM-1998-208444, 1998.

²⁹Tsieh, T., “An Investigation of Separated Flow About a Hemisphere-Cylinder at 0- to 19-Deg Incidence in the Mach Number Range from 0.6 to 1.5,” AEDC-TR-76-112, 1976.

³⁰Vassberg, J. C., DeHaan, M. A., Rivers, S. M., and Wahls, R. A., “Development of a Common Research Model for Applied CFD Validation Studies,” *26th AIAA Applied Aerodynamics Conference*, AIAA Paper 2008-6919, Honolulu, HI, 2008.

³¹Vassberg, J. C., “A Unified Baseline Grid about the Common Research Model Wing-Body for the Fifth AIAA CFD Drag Prediction Workshop,” *29th AIAA Applied Aerodynamics Conference*, AIAA Paper 2011-3508, Honolulu, HI, 2011.

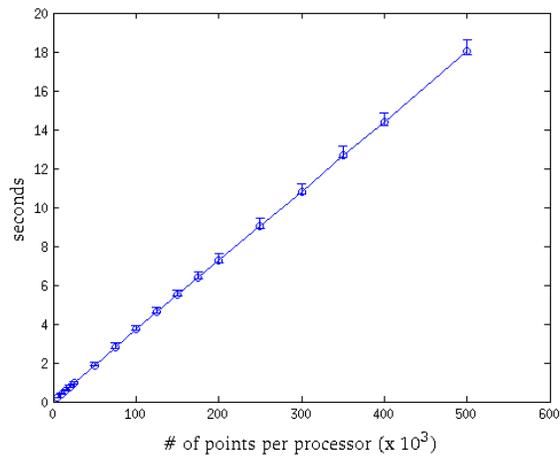


Figure 1. Variation of the TauBench work unit in seconds is shown versus the partition size (the number of grid points per processor). The computations are performed on 192 processors. 10 independent tests were performed for each partition size. The bar shows the variation of time for the same partition size. The circle marker shows the mean work unit time.

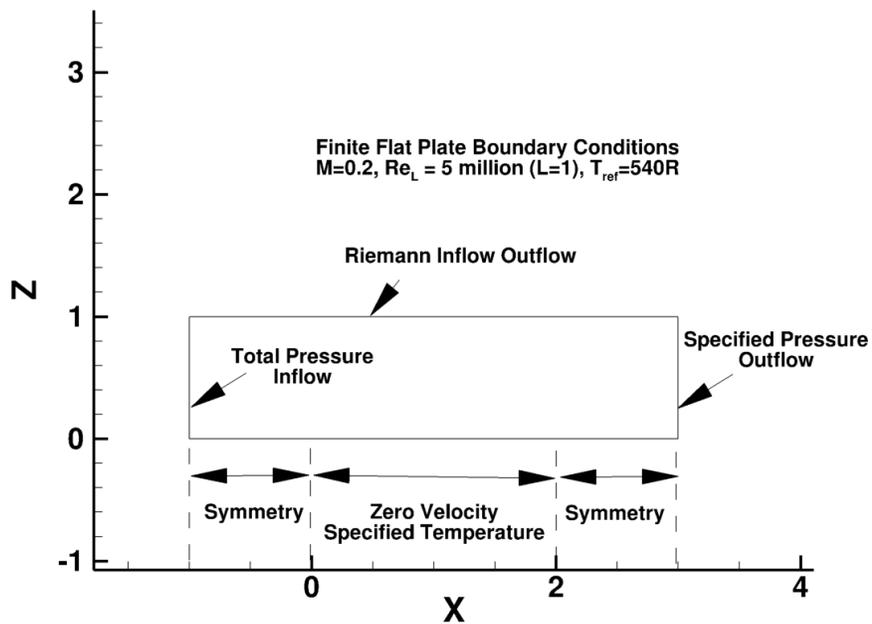
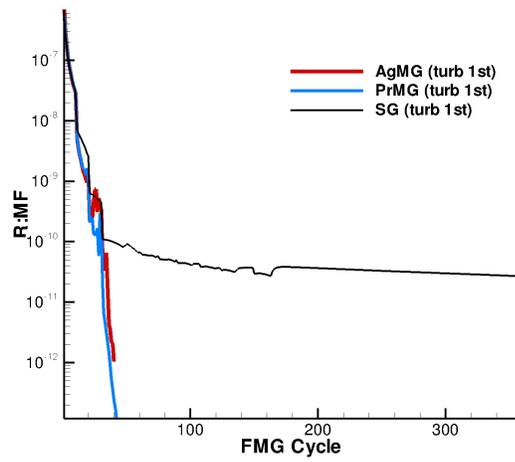
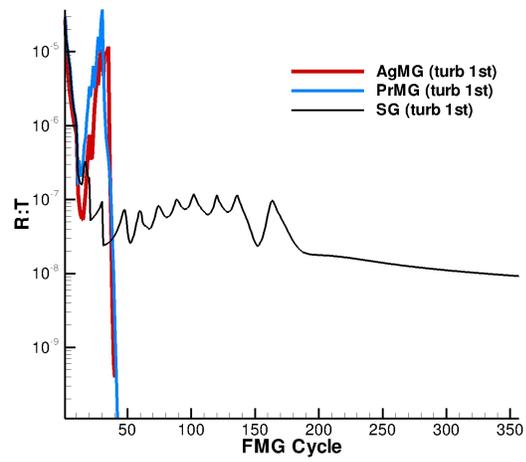


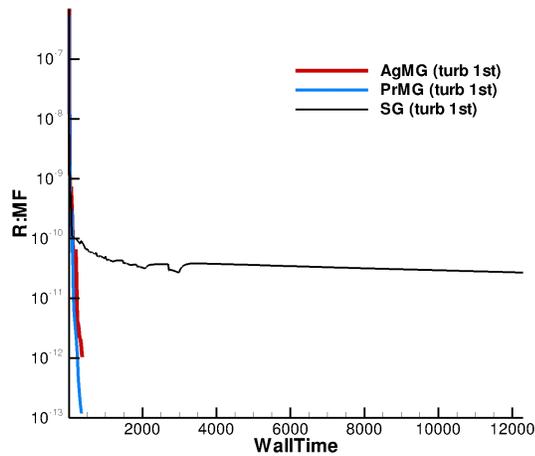
Figure 2. 2D flat plate: Geometry, flow and boundary conditions.



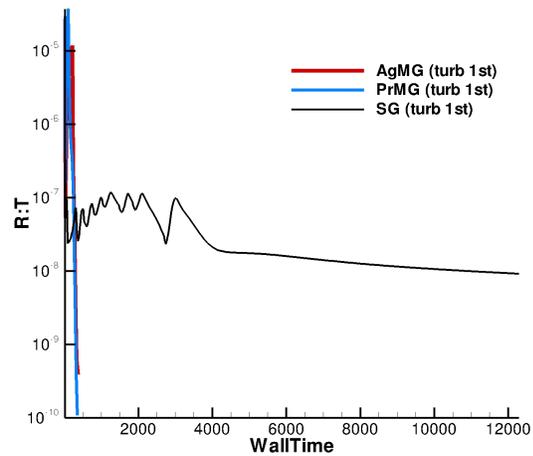
(a) Mean-flow residuals. Cycles to solution.



(b) Turbulence-model residuals. Cycles to solution.



(c) Mean-flow residuals. Time to solution.



(d) Turbulence-model residuals. Time to solution.

Figure 3. Residual convergence history for a 2D benchmark turbulent flow over a flat plate geometry from the TMR website is shown. Computations are performed on a 2561×769 quadrilateral grid, 16 processors, with a 1-st order discretization of the convection term in the SA equation. R:MF is the L_2 -norm over residuals of four mean-flow equations. R:T is the L_2 -norm of the SA turbulence-model equation residual. The wall time is measured in seconds. The AgMG and PrMG solvers use 4-level FAS(2, 2) multigrid cycles; the SG cycle is equivalent to four relaxations. Initial solution is obtained by the FMG/grid-sequencing method with 10 cycles on the coarse grids, i.e., the solvers start target-grid computations after 30 cycles.

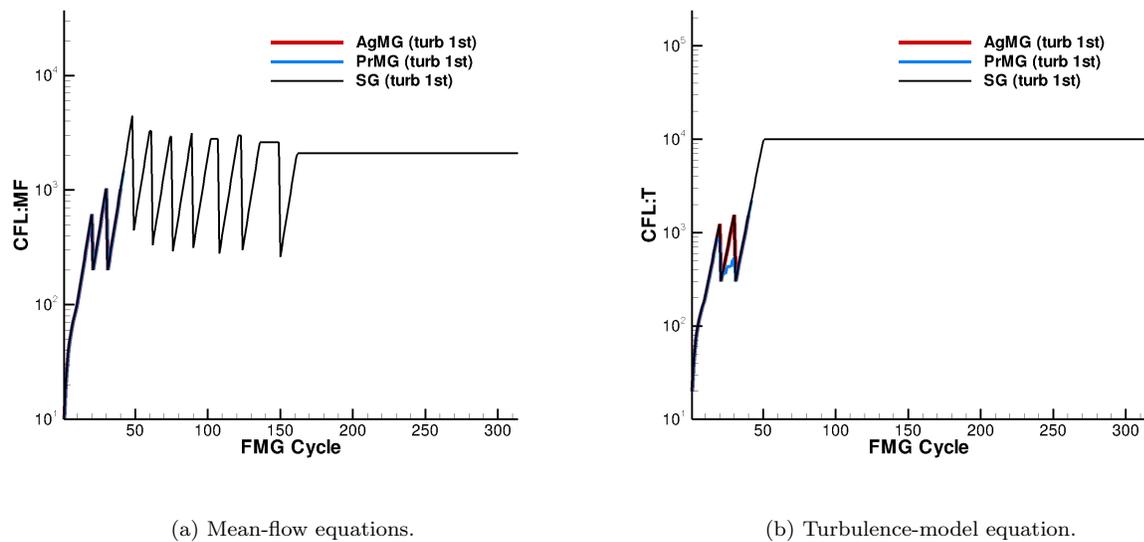


Figure 4. CFL evolution history for a 2D benchmark turbulent flow over a flat plate geometry from the TMR website is shown. Computations are performed on a 2561×769 quadrilateral grid with a 1-st order discretization of the convection term in the SA equation. The AgMG and PrMG solvers use 4-level FAS(2, 2) multigrid cycles; the SG cycle is equivalent to four relaxations, i.e., the solvers start target-grid computations after 30 cycles.

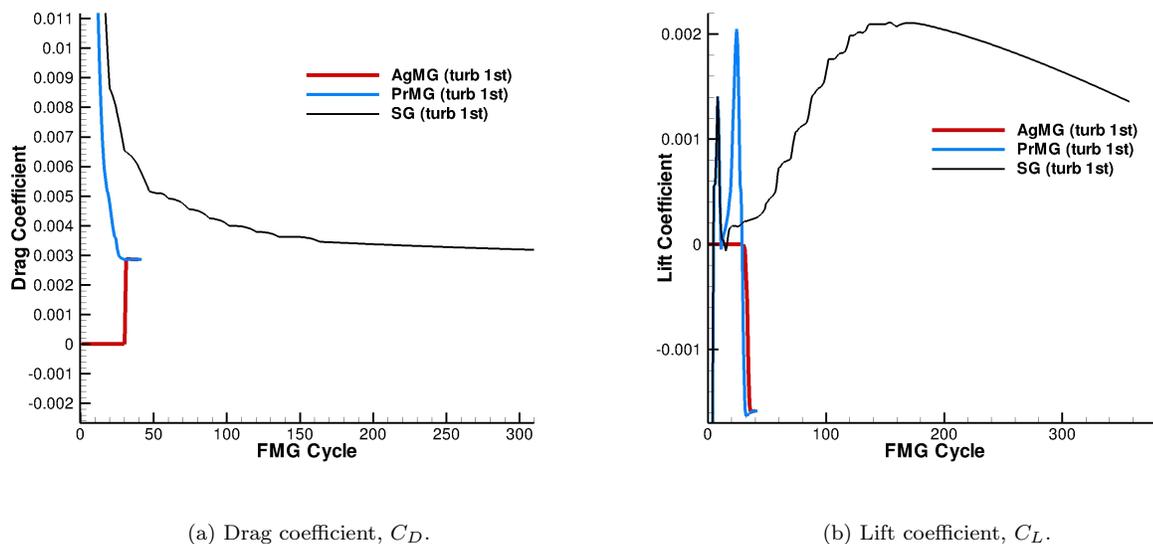


Figure 5. Convergence of lift and drag coefficients for a 2D benchmark turbulent flow over a flat plate geometry from the TMR website is shown. Computations are performed on a 2561×769 quadrilateral grid with a 1-st order discretization of the convection term in the SA equation. The AgMG and PrMG solvers use 4-level FAS(2, 2) multigrid cycles; the SG cycle is equivalent to four relaxations, i.e., the solvers start target-grid computations after 30 cycles. Only target-grid aerodynamic-coefficient evolution (i.e., evolution after 30 cycles) is shown.

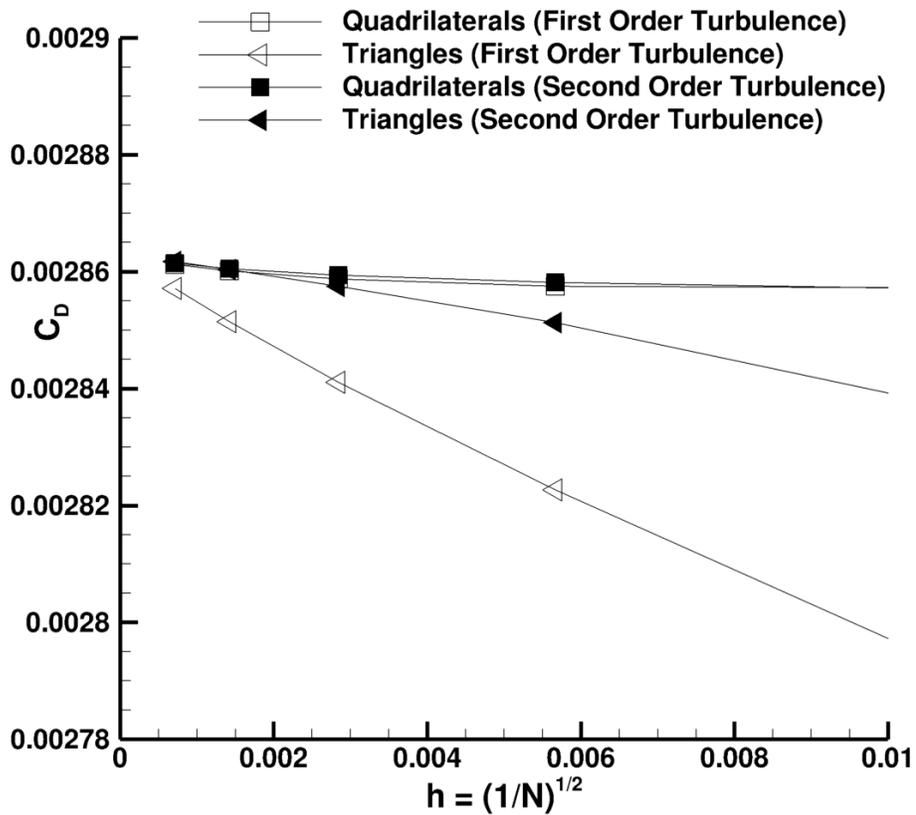


Figure 6. Grid convergence of the drag coefficient (C_D) for a 2D benchmark turbulent flow over a flat plate geometry is compared on quadrilateral and triangular grids with the first and second-order discretizations for the convection term in the turbulence-model equation. This figure is also shown in the Turbulence Model Numerical Analysis section of the TMR website.

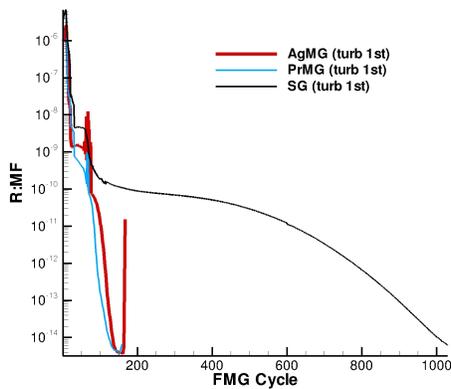
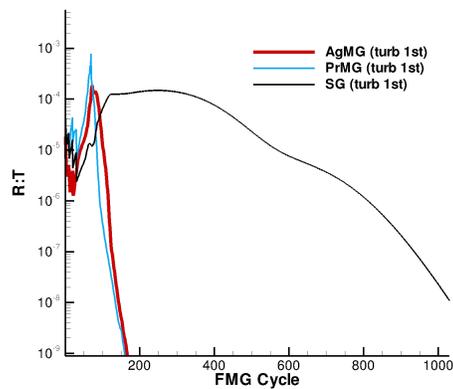
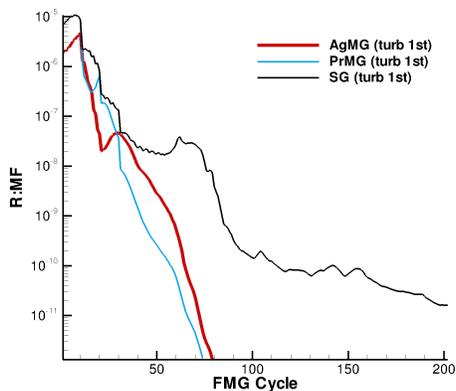
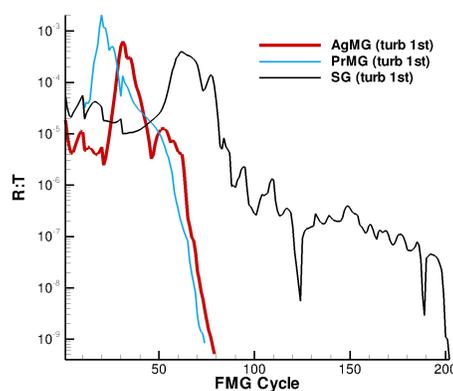
(a) Mean-flow residuals, $\alpha = 0^\circ$.(b) Turbulence-model residuals, $\alpha = 0^\circ$.(c) Mean-flow residuals, $\alpha = 10^\circ$.(d) Turbulence-model residuals, $\alpha = 10^\circ$.

Figure 7. Residual convergence history is shown versus FMG cycles for 2D benchmark turbulent flows over the NACA-0012 airfoil at two angles of attack, $\alpha = 0^\circ$ and $\alpha = 10^\circ$. These flows are validation cases from the TMR website. Computations are performed on a 1793×513 quadrilateral grid with a 1-st order discretization of the convection term in the SA equation. R:MF is the L_2 -norm over residuals of four mean-flow equations. R:T is the L_2 -norm of the SA turbulence-model equation residual. The PrMG solver use 4-level FAS(2, 2) multigrid cycle. The AgMG solver use 3-level FAS(3, 3) multigrid cycle. The SG cycle is equivalent to four relaxations. Initial solution is obtained by the FMG/grid-sequencing method with 10 cycles on the coarse grids. The PrMG and SG solvers use 4 grid levels, i.e., the solvers start target-grid computations after 30 cycles. The AgMG solver uses 3 grid levels, i.e., the solver starts target-grid computations after 20 cycles.

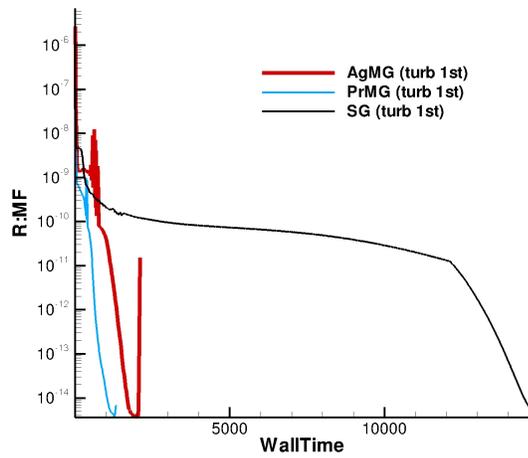
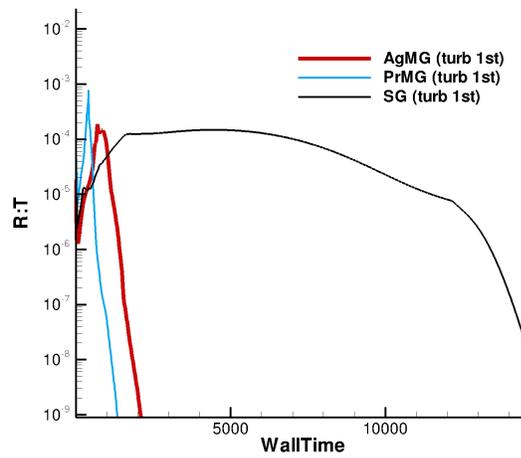
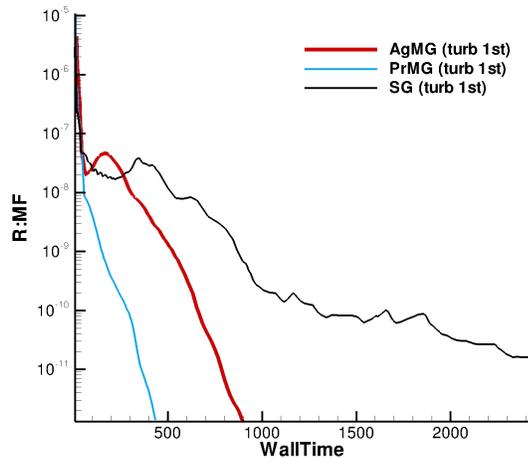
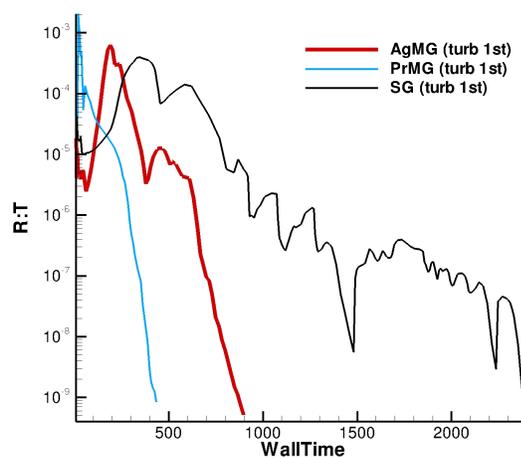
(a) Mean-flow residuals, $\alpha = 0^\circ$.(b) Turbulence-model residuals, $\alpha = 0^\circ$.(c) Mean-flow residuals, $\alpha = 10^\circ$.(d) Turbulence-model residuals, $\alpha = 10^\circ$.

Figure 8. Residual convergence history is shown versus the wall-clock time (in seconds) for 2D benchmark turbulent flows over the NACA-0012 airfoil at two angles of attack, $\alpha = 0^\circ$ and $\alpha = 10^\circ$. Computations are performed on a 1793×513 quadrilateral grid, 16 processors, with a 1-st order discretization of the convection term in the SA equation. R:MF is the L_2 -norm over residuals of four mean-flow equations. R:T is the L_2 -norm of the SA turbulence-model equation residual.

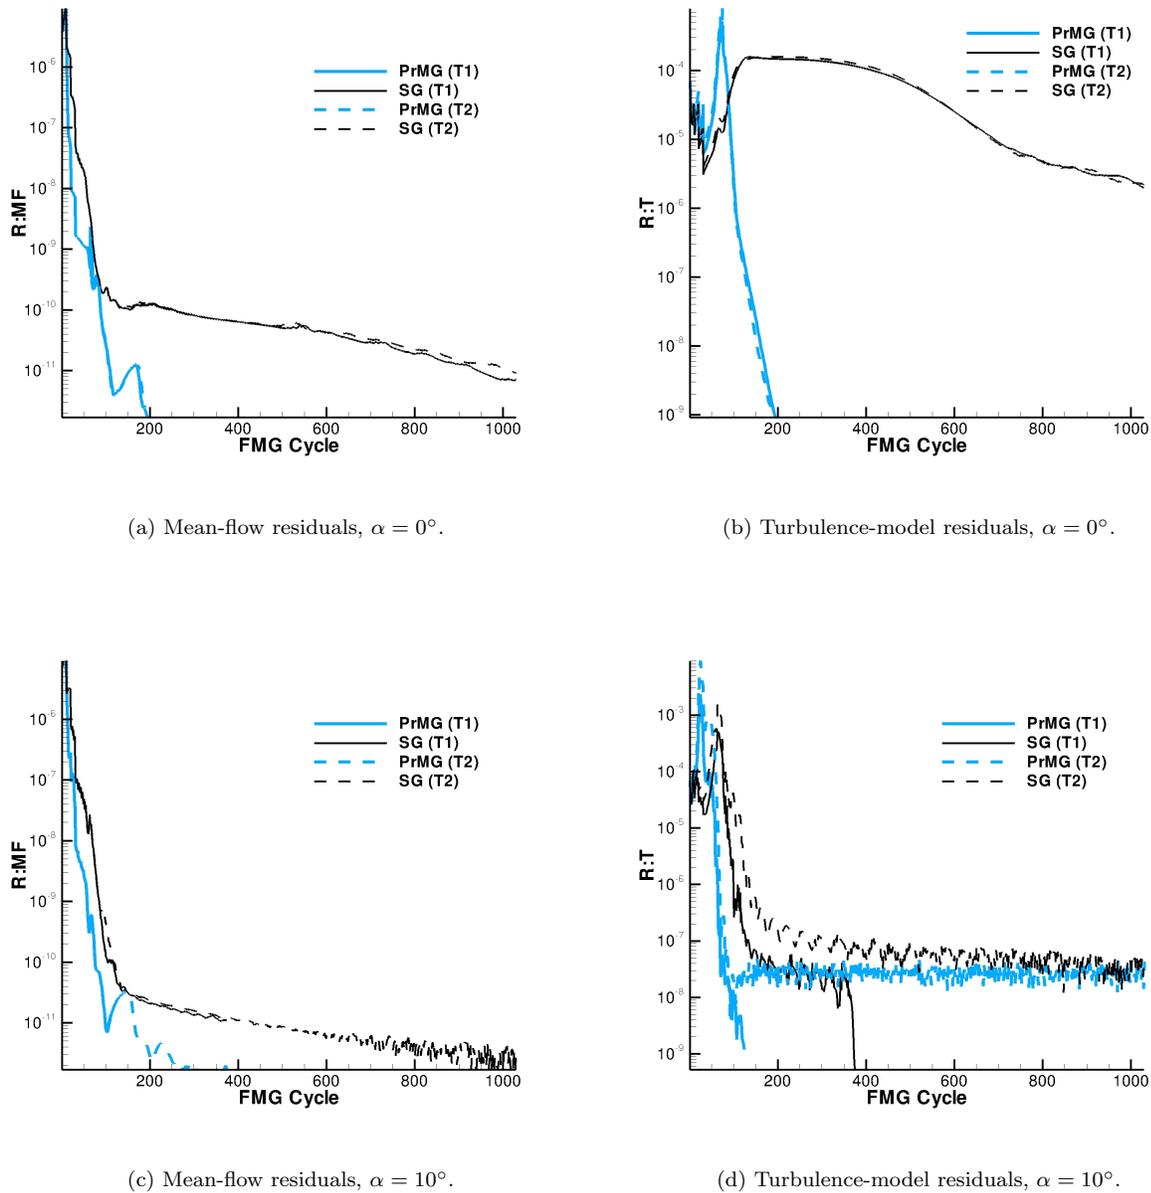
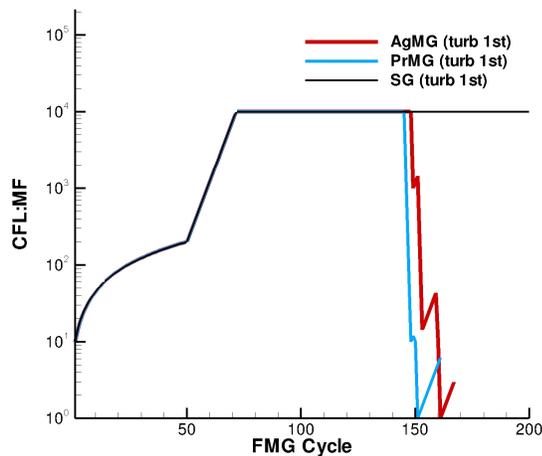
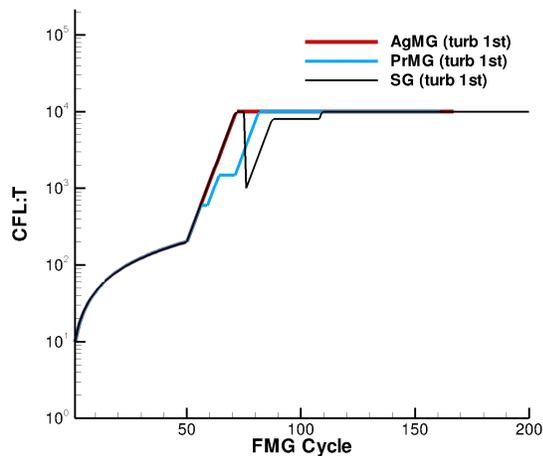


Figure 9. Residual convergence history is shown versus FMG cycles for 2D benchmark turbulent flows over the NACA-0012 airfoil at two angles of attack, $\alpha = 0^\circ$ and $\alpha = 10^\circ$. Computations are performed on a 1793×513 triangular grid. Convergence plots corresponding to 1-st and 2-nd order discretizations of the convection term in the SA equation are denoted (T1) and (T2), respectively. R:MF is the L_2 -norm of residuals of four mean-flow equations. R:T is the L_2 -norm of the SA turbulence-model equation residual. The PrMG solver uses 4-level FAS(2, 2) multigrid cycle. The SG cycle is equivalent to four relaxations. Initial solutions are obtained by the 4-level FMG/grid-sequencing method with 10 cycles on the coarse grids, i.e., the solvers start target-grid computations after 30 cycles.



(a) Mean-flow equations.



(b) Turbulence-model equation.

Figure 10. CFL evolution history is shown for a 2D benchmark turbulent flow over the NACA-0012 airfoil at $\alpha = 0^\circ$. Computations are performed on a 1793×513 quadrilateral grid, 16 processors, with a 1-st order discretization of the convection term in the SA equation.

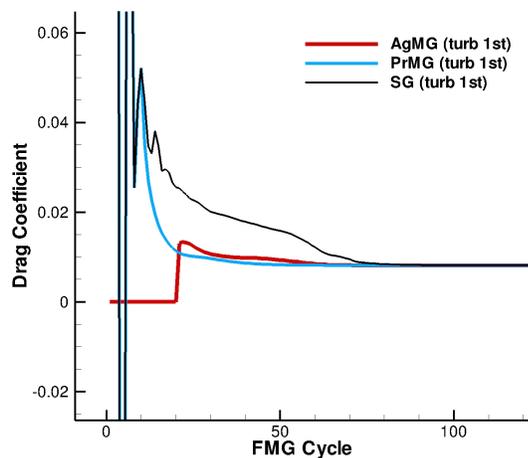
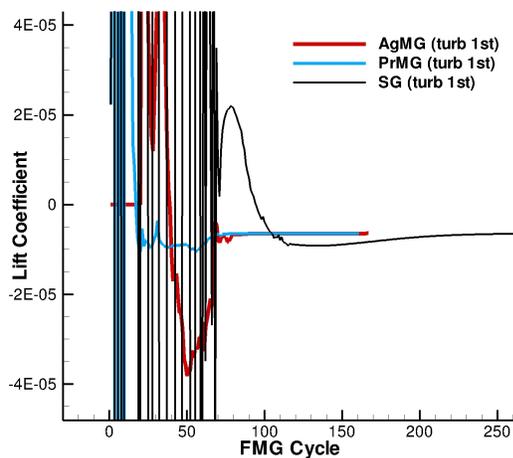
(a) Drag coefficient, C_D .(b) Lift coefficient, C_L .

Figure 11. Convergence of drag and lift coefficients is shown on a 1793×513 quadrilateral grid versus FMG cycles for a 2D benchmark turbulent flow over the NACA-0012 airfoil at $\alpha = 0^\circ$. The convection term of the SA turbulence-model equation is discretized with a first order scheme. The PrMG solver use 4-level FAS(2, 2) multigrid cycle. The AgMG solver use 3-level FAS(3, 3) multigrid cycle. The SG cycle is equivalent to four relaxations. Initial solution is obtained by the FMG/grid-sequencing method with 10 cycles on the coarse grids. The PrMG and SG solvers use 4 grid levels, i.e., the solvers start target-grid computations after 30 cycles. The AgMG solver uses 3 grid levels; only target-grid aerodynamic-coefficient evolution (i.e., evolution after 20 cycles) is shown.

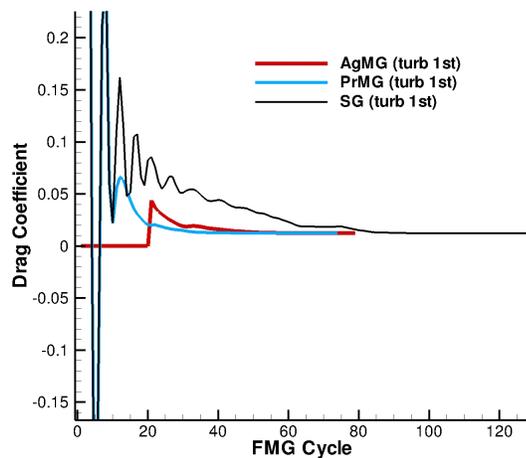
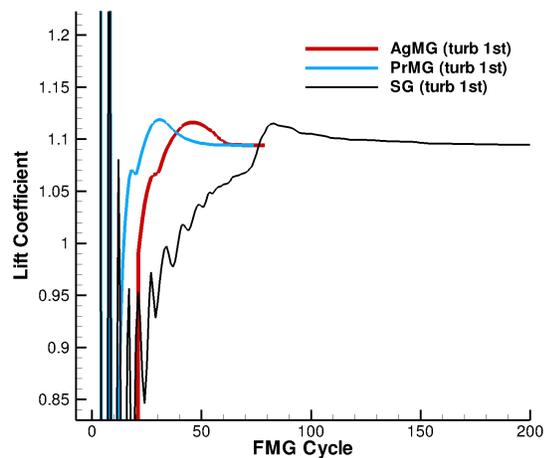
(a) Drag coefficient, C_D .(b) Lift coefficient, C_L .

Figure 12. Convergence of drag and lift coefficients is shown on a 1793×513 quadrilateral grid versus FMG cycles for a 2D benchmark turbulent flow over the NACA-0012 airfoil at $\alpha = 10^\circ$. The convection term of the SA turbulence-model equation is discretized with a first order scheme. The PrMG solver use 4-level FAS(2, 2) multigrid cycle. The AgMG solver use 3-level FAS(3, 3) multigrid cycle. The SG cycle is equivalent to four relaxations. Initial solution is obtained by the FMG/grid-sequencing method with 10 cycles on the coarse grids. The PrMG and SG solvers use 4 grid levels, i.e., the solvers start target-grid computations after 30 cycles. The AgMG solver uses 3 grid levels; only target-grid aerodynamic-coefficient evolution (i.e., evolution after 20 cycles) is shown.

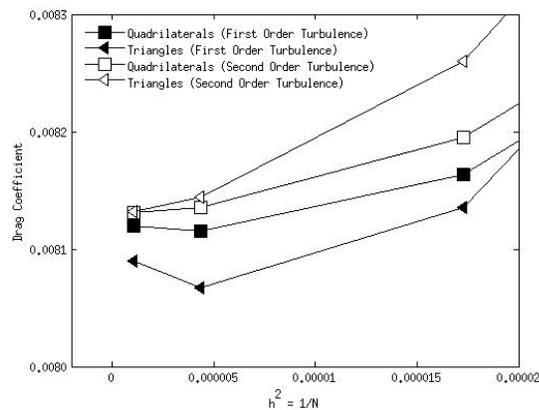
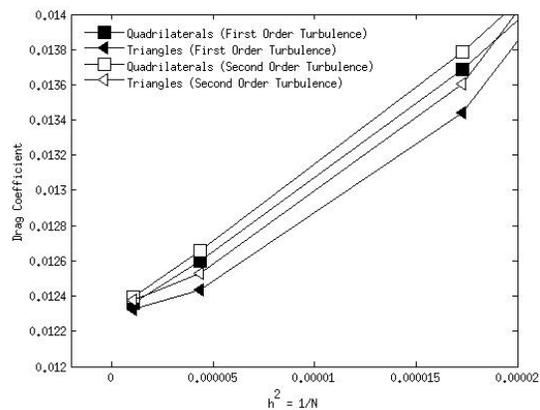
(a) $\alpha = 0^\circ$.(b) $\alpha = 10^\circ$.

Figure 13. Grid convergence of the drag coefficient for 2D benchmark turbulent flows over the NACA-0012 airfoil at two angles of attack, $\alpha = 0^\circ$ and $\alpha = 10^\circ$ is compared on quadrilateral and triangular grids with the first and second-order discretizations for the convection term in the turbulence-model equation. The family of quadrilateral grids is taken from the TMR website. The triangular grids are derived from the corresponding quadrilateral ones by dividing each quadrilateral element into two triangles. Three finest grids of the family are shown, the finest one is a 1793×513 grid.

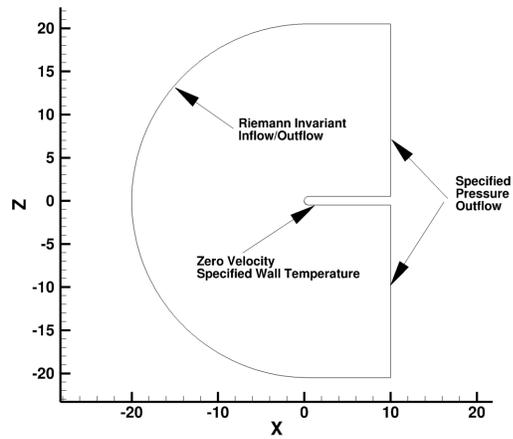


Figure 14. 3D hemisphere cylinder: Geometry and boundary conditions.

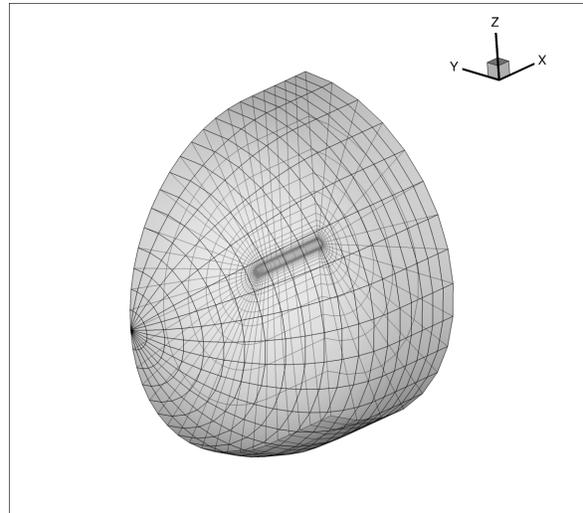


Figure 15. 3D hemisphere cylinder: View of a primal grid of Family I. Family I is based on hexahedral elements and uses polar grids over the hemisphere. Near the polar axis, the hexahedral elements degenerate into triangular prisms, leading to polar singularity. For nodes on the polar axis, the number of edge-connected neighbors increases on finer grids.

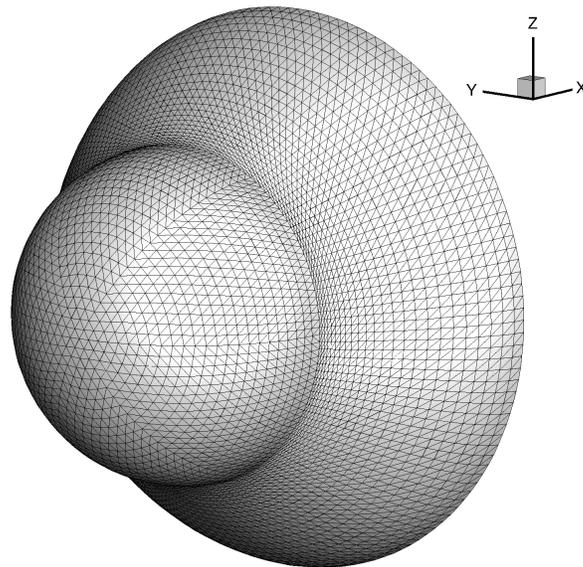
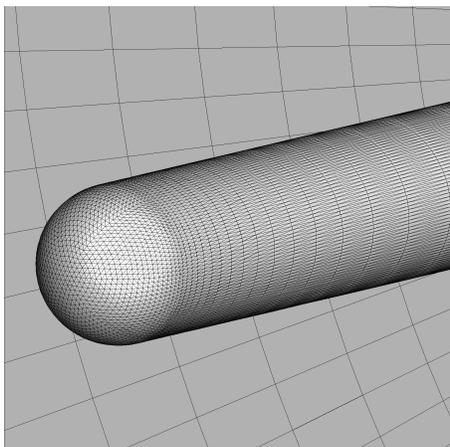
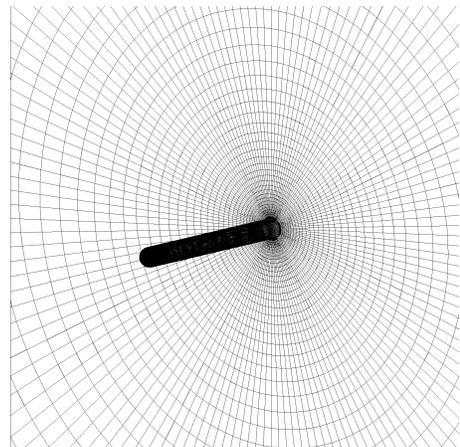


Figure 16. 3D hemisphere cylinder: Shape of the outer boundary for grids of Family II facilitates generation of relatively isotropic grids near the far-field boundary.

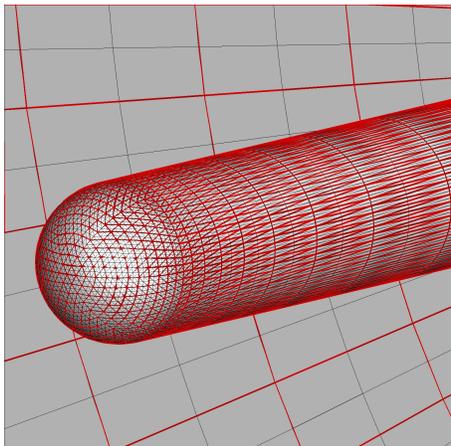


(a) Surface of a fine primal grid.

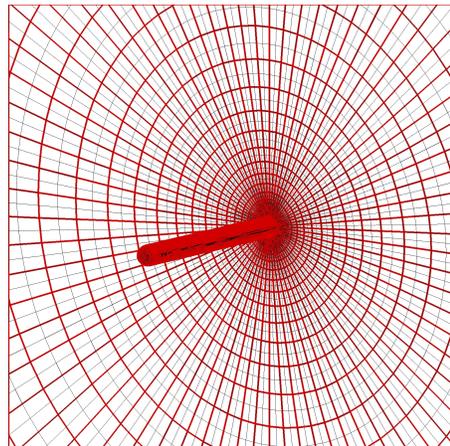


(b) Far view of a fine primal grid.

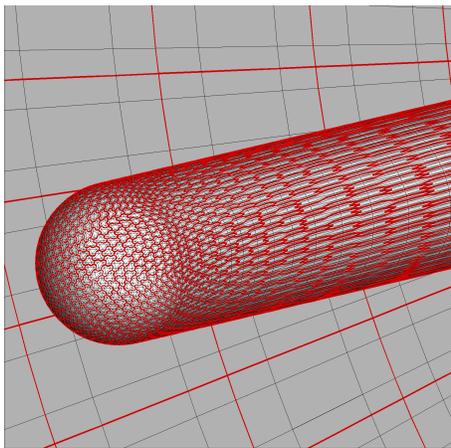
Figure 17. 3D hemisphere cylinder: Primal fine grids of Family II. Primal grids of Family II are prismatic structured grids without polar singularity, based on triangular prisms. The grid over the hemisphere is initially divided into six triangular sectors. Finer primal surface grids are generated through subdividing each triangular surface element into 4 smaller triangles by edge bisecting. Each interior surface node is guaranteed to be edge-connected to exactly six other surface nodes.



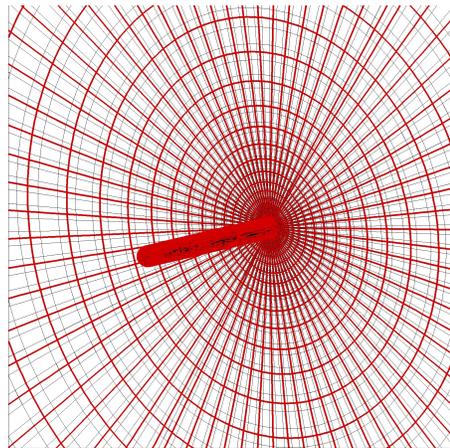
(a) Surface of a coarse primal grid.



(b) Far view of a coarse primal grid.

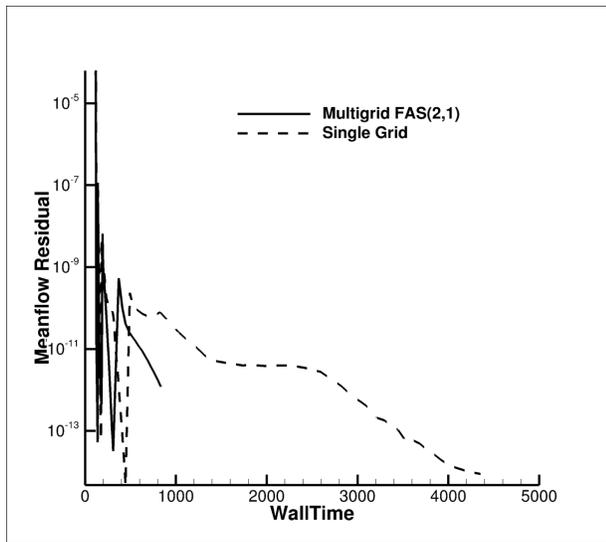


(c) Surface of an agglomerated grid.

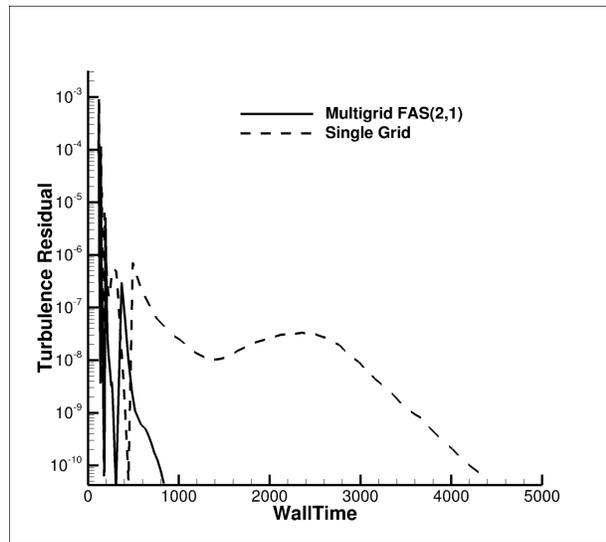


(d) Far view of an agglomerated grid.

Figure 18. 3D hemisphere cylinder: Coarse primal and agglomerated grids of Family II. The black and red lines denote fine and coarse grids, respectively.



(a) Mean-flow residuals.



(b) Turbulence-model residuals.

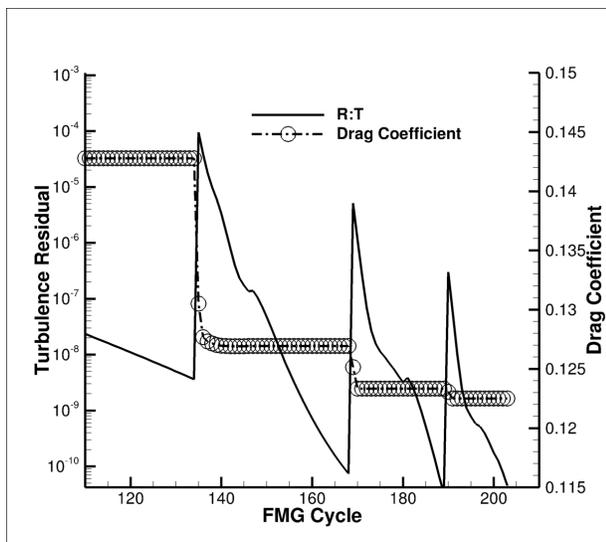
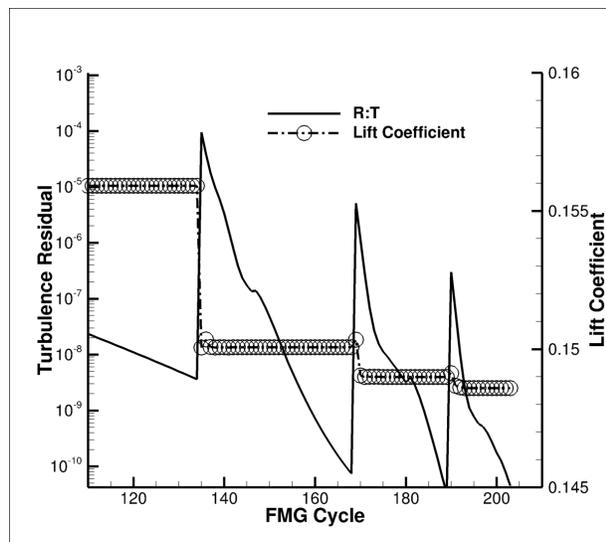
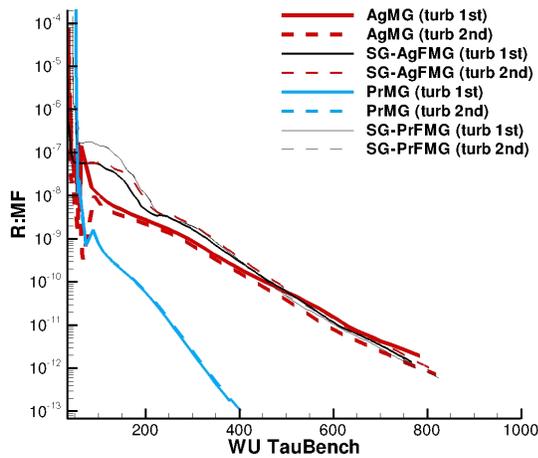
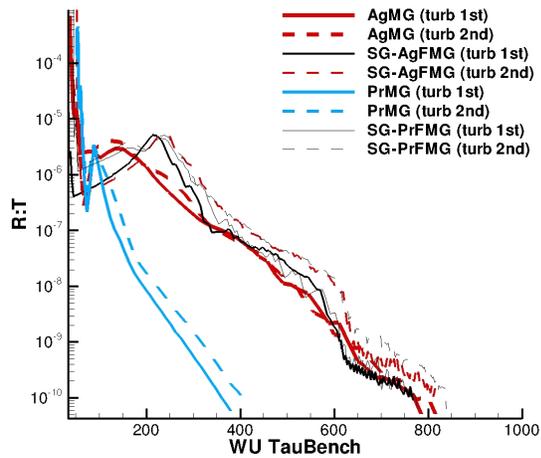
(c) Drag coefficient, C_D .(d) Lift coefficient, C_L .

Figure 19. Convergence of residuals and lift and drag coefficients is shown for a turbulent flow over the 3D Hemisphere cylinder. Computations are performed on a 20M fully hexahedral grid of Family I. R:MF is the L_2 -norm of residuals of five mean-flow equations. R:T is the L_2 -norm of the SA turbulence-model equation residual. A first-order approximation to the turbulence convection term is used. The computations are performed on 192 processors. The residual tolerance is set to be 10^{-10} . The time is the wall-clock time in seconds. In this computational environment, the TauBench work unit computed for 100K nodes/processor is about 3.674 seconds. The PrMG solver uses a 4-level FAS(2, 1) cycle. An FMG method is used to form initial approximations for multigrid cycles. FMG solutions converge to the required tolerance on all grids. The turbulence-equation residuals are also shown with plots for the drag and lift coefficients to distinguish between forces computed on different grids.



(a) Mean-flow residuals.



(b) Turbulence-model residuals.

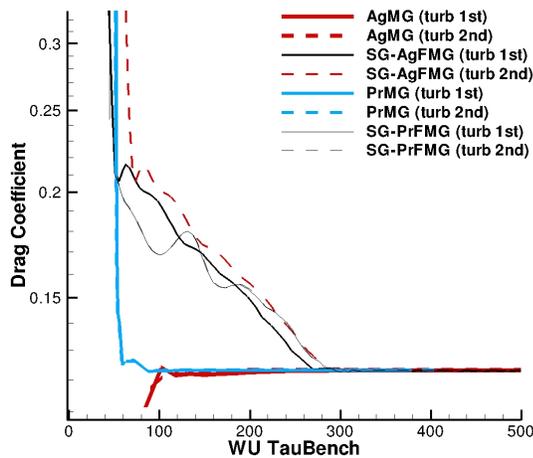
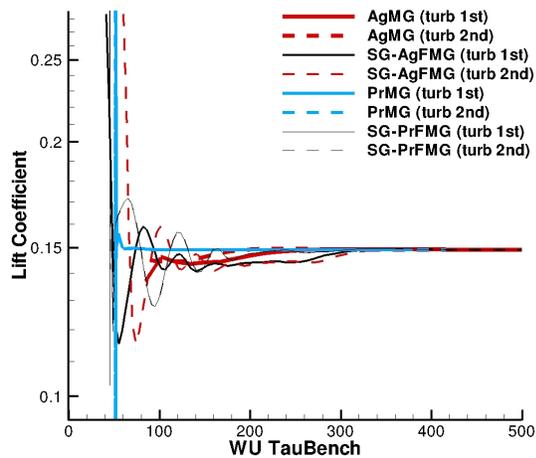
(c) Drag coefficient, C_D .(d) Lift coefficient, C_L .

Figure 20. Convergence of residuals and lift and drag coefficients is shown for a turbulent flow over the 3D Hemisphere cylinder. Computations are performed on a 20M fully prismatic grid of Family II. R:MF is the L_2 -norm of residuals of five mean-flow equations. R:T is the L_2 -norm of the SA turbulence-model equation residual. The first and second-order approximations to the turbulence convection term are used. The computations are performed on 192 processors. The residual tolerance is set to be 10^{-10} . The time is measured in TauBench work units. In this computational environment, the TauBench work unit computed for 100K nodes/processor is about 3.674 seconds. The PrMG and AgMG solvers use 4-level FAS(2, 2) and FAS(3, 3) cycles, respectively. An FMG method with 10 coarse-grid cycles is used to form initial approximations for multigrid cycles. The SG-AgFMG and SG-PrFMG notations indicate that the SG solver uses grid sequencing approach based either on agglomerated or primal grids, respectively.

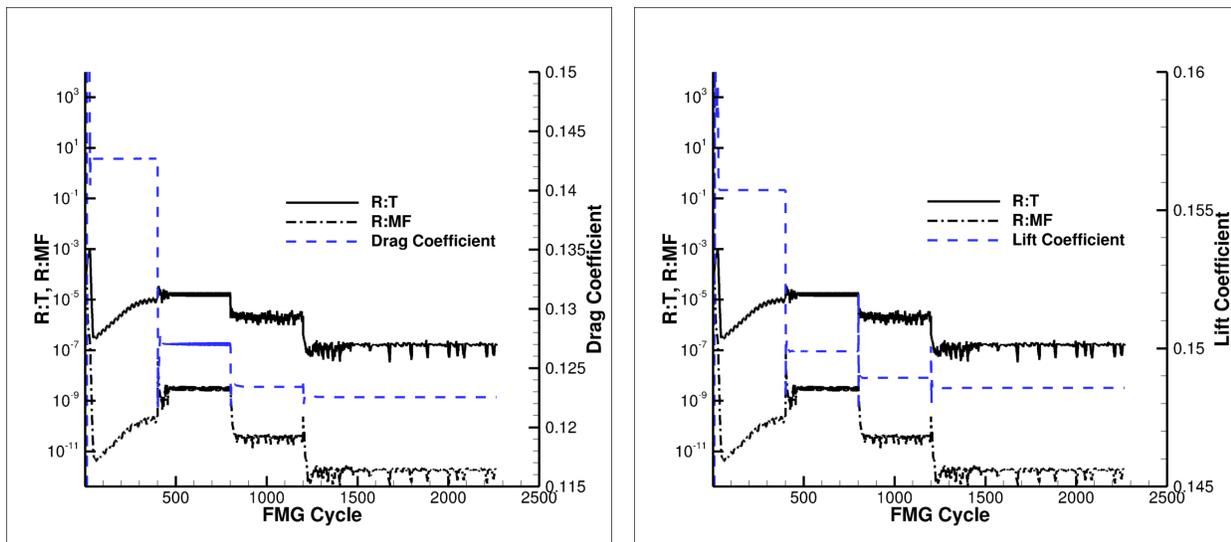
(a) Drag coefficient, C_D .(b) Lift coefficient, C_L .

Figure 21. Convergence of residuals and lift and drag coefficients is shown for a turbulent flow over the 3D Hemisphere cylinder. Computations are performed on a 20M fully hexahedral grid of Family I. R:MF is the L_2 -norm of residuals of five mean-flow equations. R:T is the L_2 -norm of the SA turbulence-model equation residual. A second-order approximation to the turbulence convection term is used. The PrMG solver uses a 4-level FAS(2, 1) cycle. The residuals are shown with plots for the drag and lift coefficients to distinguish between forces computed on different grids.

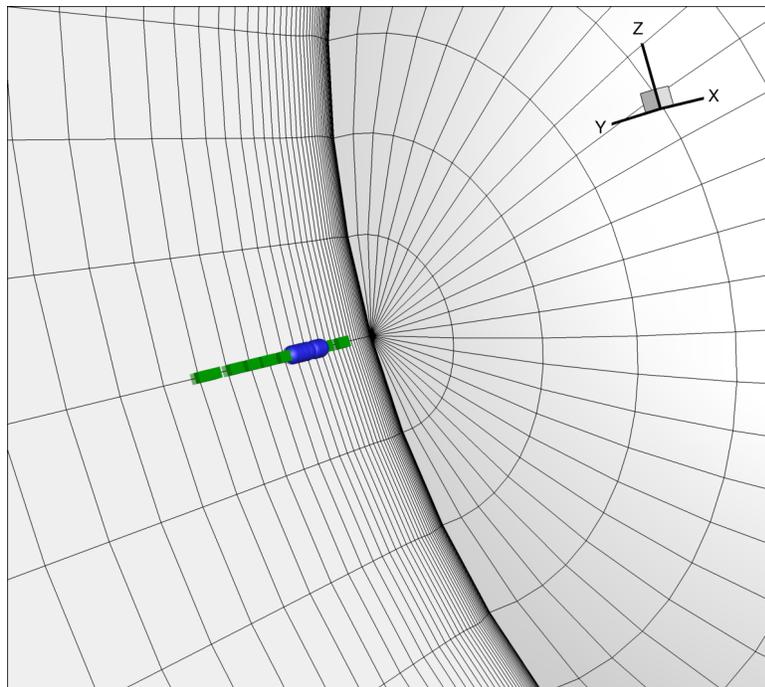
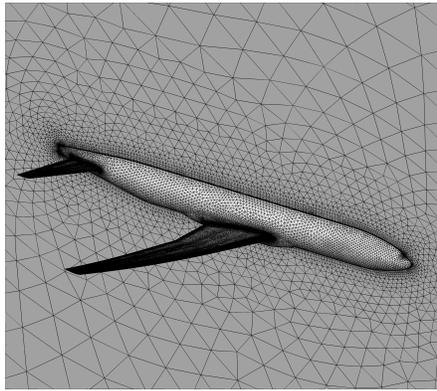
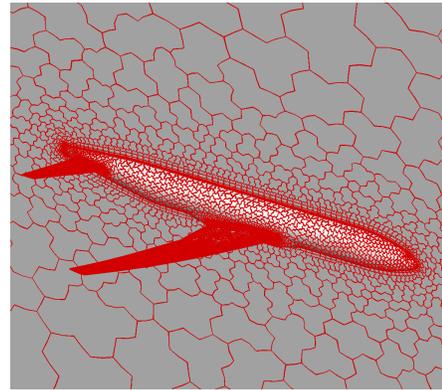


Figure 22. Location of the largest residuals and the negative turbulence variables is illustrated for a grid of Family I. Blue spheres illustrate locations of the largest residuals over last 20 cycles on the finest 20M grid. Green cubes show the locations of negative turbulence working variables in the last cycle on this grid. The background grid shown is the coarsest one in Family I.

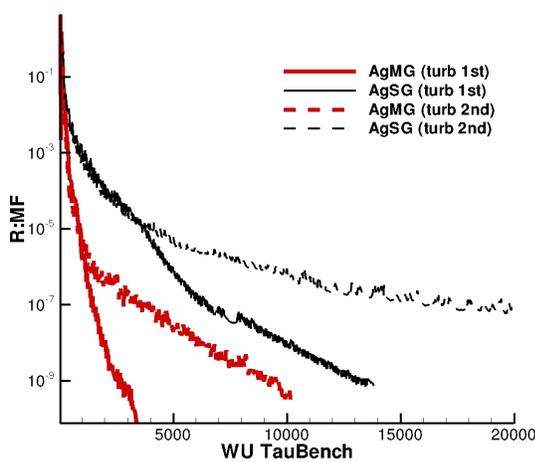


(a) Surface of a fine primal grid.

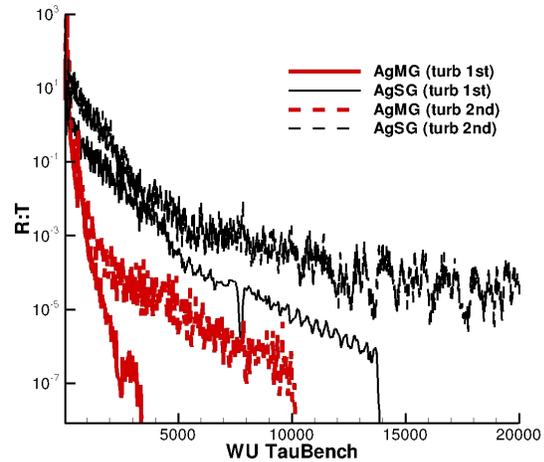


(b) Surface of an agglomerated grid.

Figure 23. DPW-4 configuration: Fully unstructured mixed-element primal and corresponding agglomerated grids.



(a) Mean-flow residuals.



(b) Turbulence-model residuals.

Figure 24. Time to convergence is shown for a turbulent flow over the DPW-4 wing-body-tail configuration. Computations are performed on a (medium) 10M mixed-element grid. R:MF is the L_2 -norm of residuals of five mean-flow equations. R:T is the L_2 -norm of the SA turbulence-model equation residual. The first and second-order approximations to the turbulence convection term are used. The computations are performed on 180 processors. The residual tolerance is set to be 10^{-8} . The time is measured in TauBench work units. In this computational environment, the TauBench work unit computed for 57K nodes/processor is about 2.473 seconds. The AgMG solver uses 3-level FAS(3, 3) cycle. An FMG method with 10 coarse-grid cycles is used to form initial approximations for multigrid cycles. The AgSG notation indicates that the SG solver uses grid sequencing approach based on agglomerated grids.

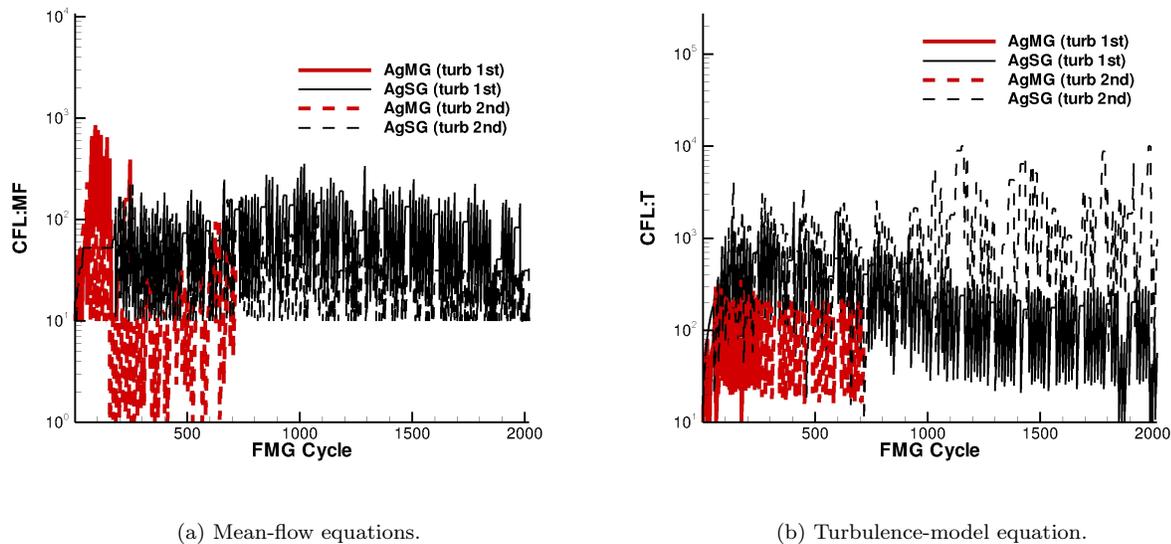


Figure 25. CFL variation is shown for a turbulent flow over the DPW-4 wing-body-tail configuration. Computations are performed on a (medium) 10M mixed-element grid generated for DPW-4. CFL:MF is the CFL number for the mean-flow equations. CFL:T is the CFL number for the turbulence-model equation. The first and second-order approximations to the turbulence convection term are used. An FMG method with 10 coarse-grid cycles is used to form initial approximations for multigrid cycles. The AgSG notation indicates that the SG solver uses grid sequencing approach based on agglomerated grids. The AgMG solver uses 3-level FAS(3, 3) cycle. The AgSG cycle is equivalent to six iterations.

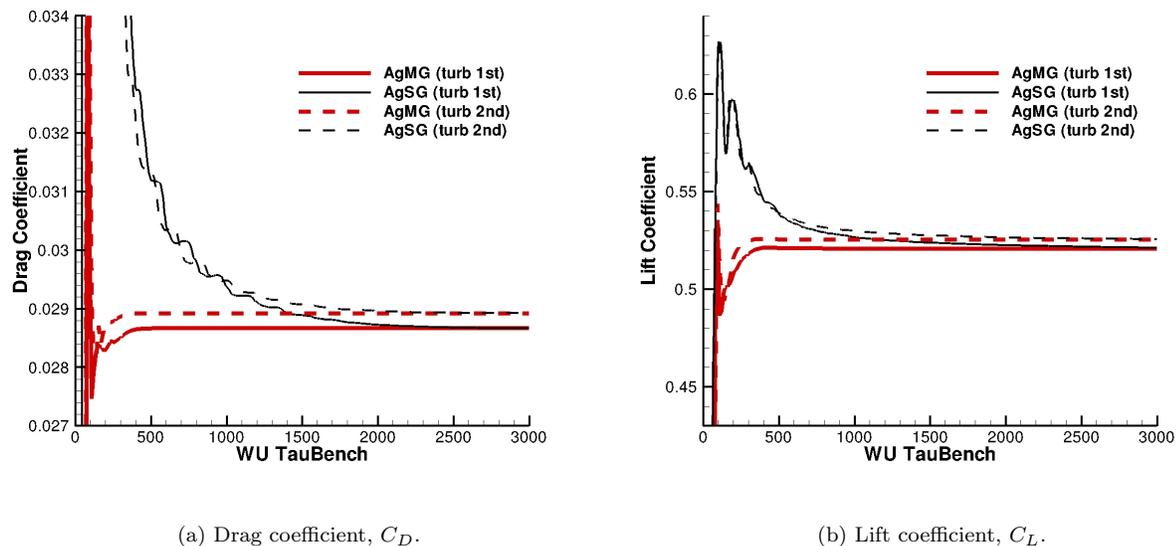
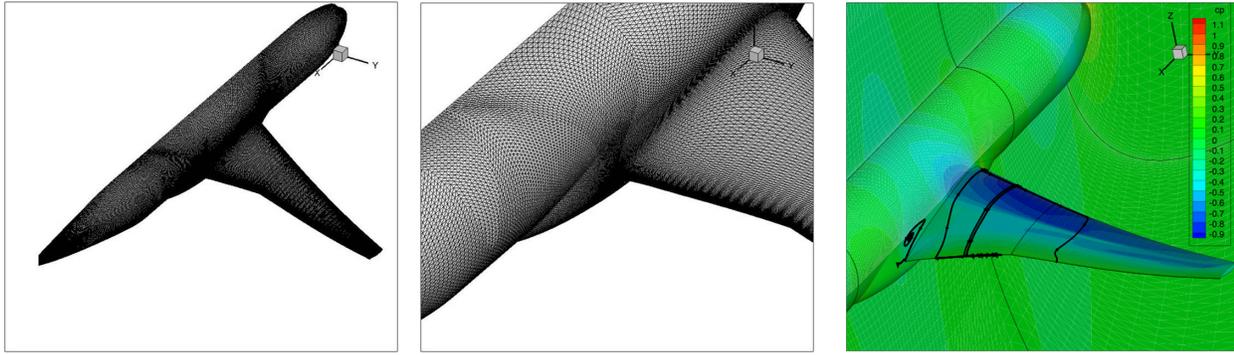


Figure 26. Convergence of lift and drag coefficients is shown for a turbulent flow over the DPW-4 wing-body-tail configuration. Computations are performed on a (medium) 10M mixed-element grid on 180 processors. The time is measured in TauBench work units. In this computational environment, the TauBench work unit computed for 56K nodes/processor is about 2.473 seconds. The first and second-order approximations to the turbulence convection term are used. The AgSG notation indicates that the SG solver uses grid sequencing approach based on agglomerated grids.

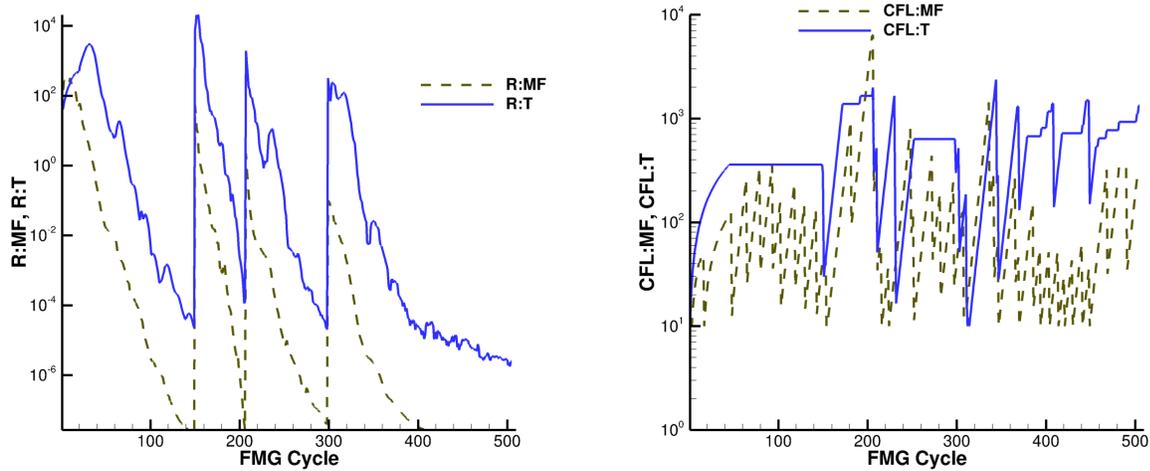


(a) Surface grid, far view.

(b) Surface grid, close view.

(c) Volume grid, pressure contours and streamlines.

Figure 27. A primal grid is shown for a DPW-5 wing-body configuration. Surface grid is triangular. Prismatic elements extend from the surface outward to include a fixed percentage of the points in the normal direction and tetrahedra fill the remainder of the domain. Pressure contours and streamlines are shown for a solution corresponding to $\alpha = 2.25^\circ$ angle of attack.



(a) Residual convergence.

(b) CFL variation.

Figure 28. Convergence histories of the residuals and CFL number in computations with the PrMG solver are shown for a DPW-5 wing-body configuration on the third coarsest workshop grid with approximately $5.2M$ grid points. R:MF is the L_2 -norm of residuals of five mean-flow equations. R:T is the L_2 -norm of the SA turbulence-model equation residual. CFL:MF is the CFL number for the mean-flow equations. CFL:T is the CFL number for the turbulence-model equation. The first-order approximation to the turbulence convection term is used. A FAS(2, 2) cycle is used. An FMG method is used to form initial approximations for multigrid cycles. The tolerance on each grid requires 8 orders of magnitude residual reduction from the maximum residual norm exhibited in the iterations on the current grid.

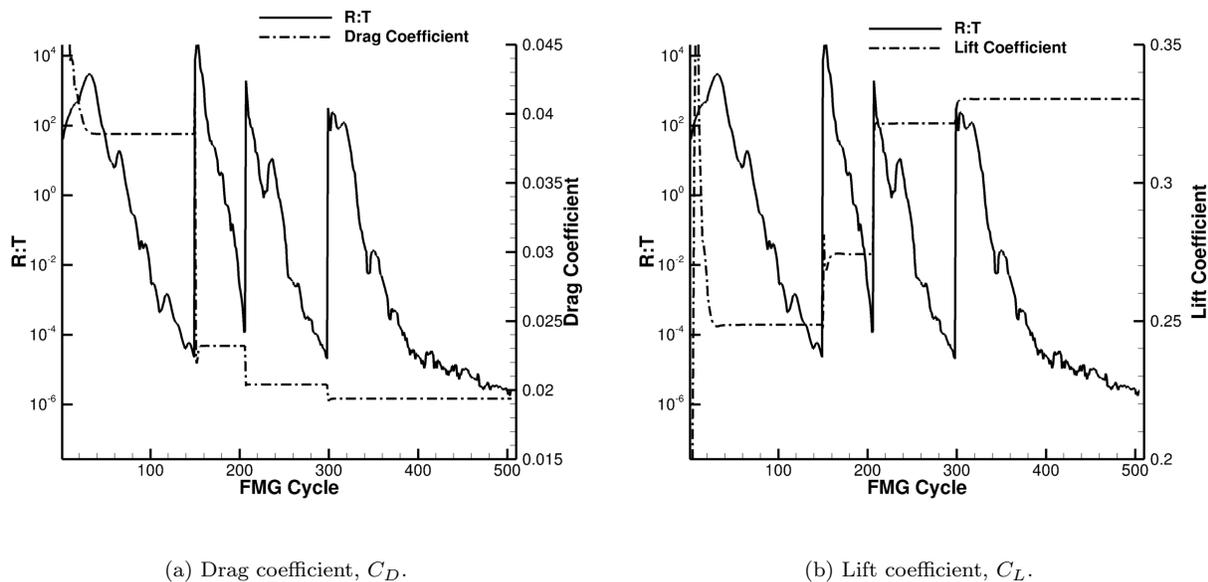


Figure 29. Convergence of lift and drag coefficients is shown for a turbulent flow over the DPW-5 wing-body configuration versus cycles. The first-order approximation to the turbulence convection term is used. The PrMG solver uses FAS(2, 2) cycled and the FMG method to form an initial approximation on fine grids. The turbulence-equation residuals are also shown to distinguish between forces computed on different grids.

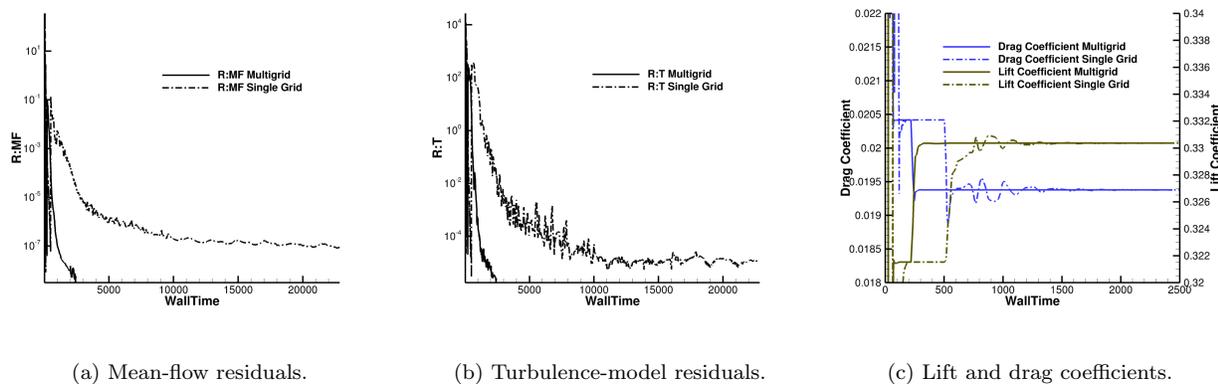


Figure 30. Comparison of the PrMG and SG solutions for a turbulent flow over the DPW-5 wing-body configuration. The computations are performed on the third coarsest workshop grid with approximately 5.2M grid points. A FAS(2, 2) cycle is used by the PrMG solver. An FMG method is used to form initial approximations for both the PrMG and SG solutions. The termination tolerance requires 8 orders of magnitude residual reduction from the maximum residual norm exhibited in the iterations. The computations are performed on 192 processors, roughly corresponding to 27K nodes per processor. The corresponding TauBench work unit is about 1 second.

Appendix A: Grid Refinement Study for NACA-0012 Airfoil

A grid refinement study for a turbulent flow around the NACA-0012 airfoil at $\alpha = 10^\circ$ angle of attack is presented. The flow is essentially incompressible (Mach number = 0.15). All solutions are obtained with the multigrid technology evaluated in the main body of the paper that dramatically improves efficiency of turbulent-flow solutions. The original intent of this study was to compare accuracy of different discretizations, analogously to the flat plate case (see Figure 6). However, to assess discretization accuracy, some information about the exact solution to the problem is required. Drag and lift coefficients computed on a family of grids are reported on the TMR website¹² for the SA turbulence model, but it is impossible to discern the values for these coefficients that would correspond to the grid-refinement limit. The results for pitching moment coefficients are not reported. The computational results have been collected from seven different codes — six codes computed on the same second finest 897×257 grid and one code applied grid adaptation. The results show a spread of about 20 drag counts (4% difference) and about 1% difference in lift computations.

The grid-refinement study was initially conducted on the four finest grids available on the website. The surface grids in this study are characterized by three parameters: the chord-wise mesh spacing at the trailing edge, the chord-wise mesh spacing at the leading edge, and the number of grid points on the surface. The finest 1793×513 grid in the family has the chord-wise mesh spacings 5×10^{-4} and 5×10^{-5} at the trailing and leading edges, respectively. The grids in this family are referenced as $1M(1)$ grids. Here $1M$ refers to a rounded number of nodes on the finest grid in the family. The number in the parenthesis is introduced for further use and indicates a reduction factor for the chord-wise mesh spacings at the trailing edge relative to the original ($1M(1)$) family at the TMR website. For example, the finest grid in the $1M(10)$ family would have the same 1793×513 dimensions, but the spacing at the trailing edge would be 5×10^{-5} . The coarser grids in a family are derived from the finest grid by successive removal of each other grid line in each dimension. The coarse grids are identified by the number after the dot: $1M(1).1$ refers to the finest grid, $1M(1).2$ refers to the next coarser grid, etc.

Initially, solutions with the residual tolerance of 10^{-9} have been computed on four $1M(1)$ grids by the PrMG solver on 16 processors with a turnaround time of 338 seconds. Surprisingly, the grid-refined values of the lift and pitching moment coefficients were more problematic to use for infinite-grid solution extrapolation than the drag coefficient values on the corresponding grids. The force and pitching moment coefficients computed on the quadrilateral grids with the first-order accurate approximation for the convection term of the turbulence-model equation are shown in Figure A.1. The collective lift and drag values from the website are represented in Figures A.1 (a) and (b) as intervals positioned at the corresponding values computed on the $1M(1).2$ grid. The lift values increase with grid refinement and then decrease on the finest grid. The finest-grid lift coefficient is on the lower fringe of the collective results shown on the TMR website. The pitching moment coefficients are continually increasing with grid refinement, but show an order property considerably less than second order.

In pursuit of a better quantification of the grid-refined aerodynamic quantities, another family of structured quadrilateral grids has been generated; the finest 7169×2049 grid has about $16M$ nodes and four times smaller chord-wise mesh spacings than the $1M(1).1$ grid. This family is referenced as $16M(1)$ family. The force and pitching moment coefficients computed on the $16M(1)$ grids are also shown in Figure A.1 and indicate the departure from the collective behavior of the TMR solutions as the grids are refined, but the values of the aerodynamic coefficients in the grid-refined limit are yet unclear. The PrMG solver on $16M(1).1$ grid uses the FMG method with 10 FAS(2, 2) cycles on coarse grids and converges residuals to the tolerance of 10^{-9} in 20 cycles. The solution has been obtained in 1202 seconds on 192 processors.

A detailed inspection of the solutions on the surface of the airfoil revealed that the pressure distribution at the trailing edge is not sufficiently resolved. Figures A.2 and A.3 illustrate the pressure and skin friction distribution on the surface. The convergence of the pressure coefficients clearly lacks any order property near the trailing edge, while the skin friction converges reasonably well. This observation motivated generation of a new grid family, $1M(10)$, that has the same dimensions as the original $1M(1)$ grids on the TMR website, but increased grid density near the trailing edge. Figure A.4 illustrates the differences between the grid families in the trailing edge area. Much improved convergence of the force and pitching moment coefficients on the $1M(10)$ grids is shown in Figure A.1. The aerodynamic coefficients computed on the $1M(10)$ grids are very close to the corresponding coefficients computed on a family of finer grids, $16M(10)$. Also the results are shown for computations on the $16M(10)$ grids with a second-order approximation for the convection term in the SA turbulence-model equation. Figures A.5 and A.6 show distributions of the pressure and skin friction on the surface near the leading and trailing edges and at an intermediate location, indicating

second-order convergence in these three locations.

Figures A.7 illustrate the convergence variation of the aerodynamic coefficients as a function of the mesh spacing at the trailing edge. Decreasing the surface mesh spacing (increasing the mesh density) near the trailing edge appears beneficial for accuracy of all aerodynamic coefficients. The drag convergence is least sensitive to the trailing edge spacing variation. The lift coefficient shows a significantly better order property for a very high grid density at the trailing edge corresponding to the $1M(40)$ grid family, but the lift coefficients computed on coarse grids are more accurate with a moderate density provided by the $1M(2)$ grid family. The pitching moment coefficient is both most accurate and demonstrates the best order property (a straight-line convergence plot) on grids of the $1M(10)$ family.

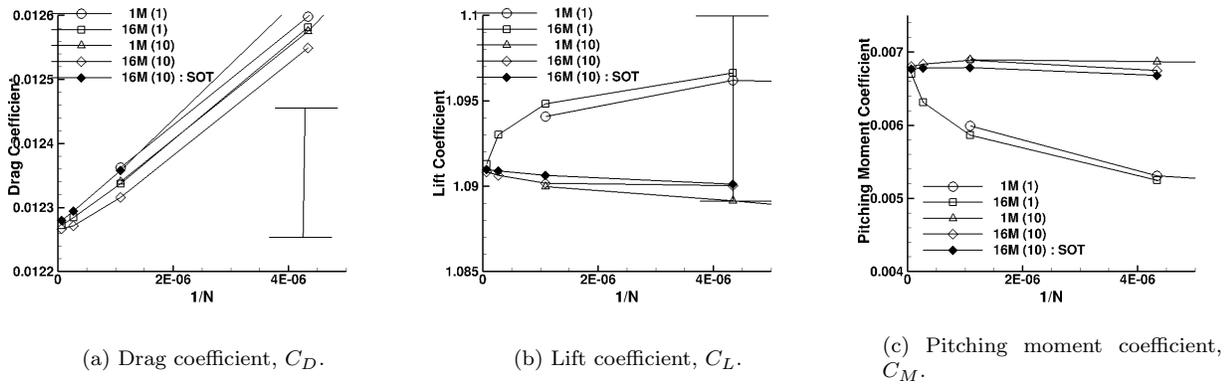
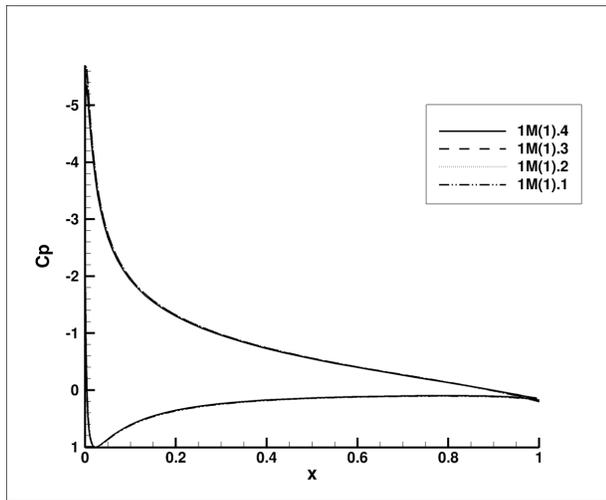
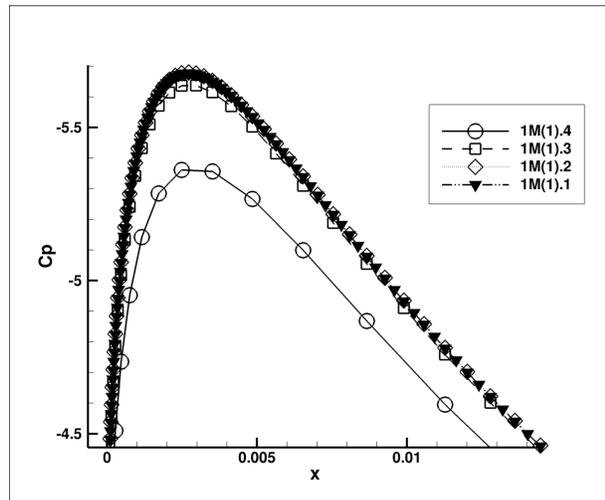


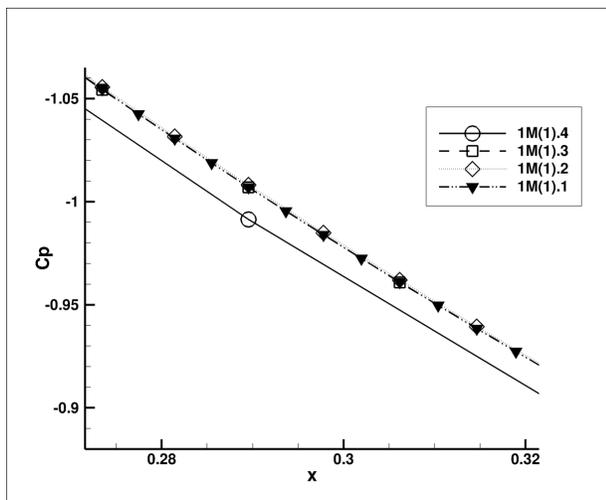
Figure A.1. The force and pitching moment coefficients are computed on the quadrilateral grids for a turbulent flow over NACA-0012 airfoil at the angle of attack of $\alpha = 10^\circ$. Abbreviation “SOT” denotes computations with a second-order approximation for the convective term in the turbulence-model equation; the first-order approximation is default. 1M(1) is the family of grids from the TMR website; the finest grid has the 1793×513 dimensions and the chord-wise mesh spacings 5×10^{-4} and 5×10^{-5} at the trailing and leading edges, respectively. 16M(1) is a family of finer grids; the finest grid has the 7169×2049 dimensions and the chord-wise mesh spacings 1.25×10^{-4} and 1.25×10^{-5} at the trailing and leading edges, respectively. 1M(10) is a family of grids with an increased mesh density at the trailing edge; the finest grid has the 1793×513 dimensions and the chord-wise mesh spacings 5×10^{-5} at both the trailing and leading edges. 16M(10) is a family of grids with an increased mesh density at the trailing edge; the finest grid has the 7169×2049 dimensions and the chord-wise mesh spacings 1.25×10^{-5} at both the trailing and leading edges. The collective lift and drag values from the TMR website are represented as intervals positioned at the values computed on the second finest 1M(1).2 grid.



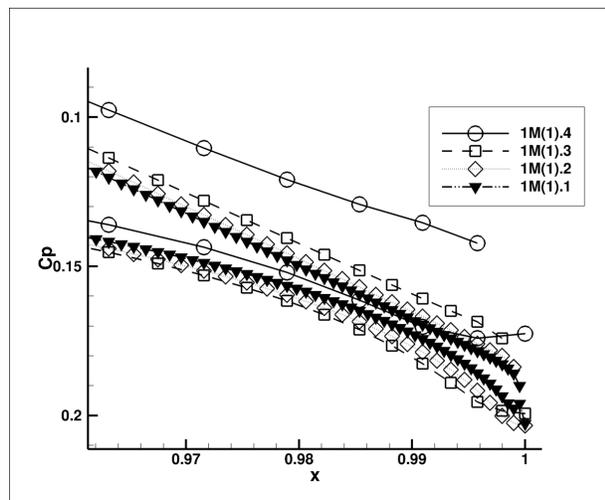
(a) Global view.



(b) Leading edge.

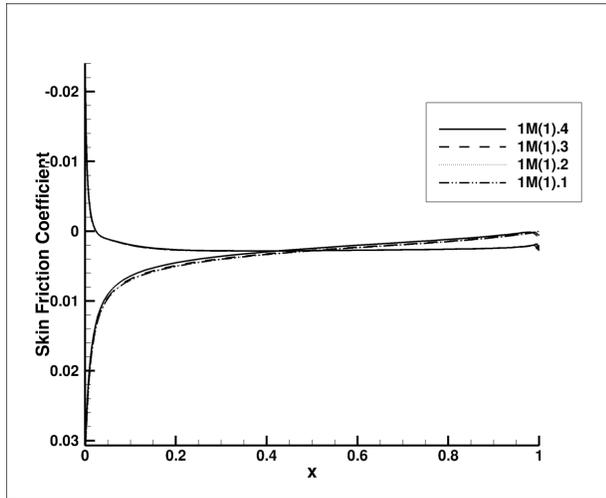


(c) Intermediate location.

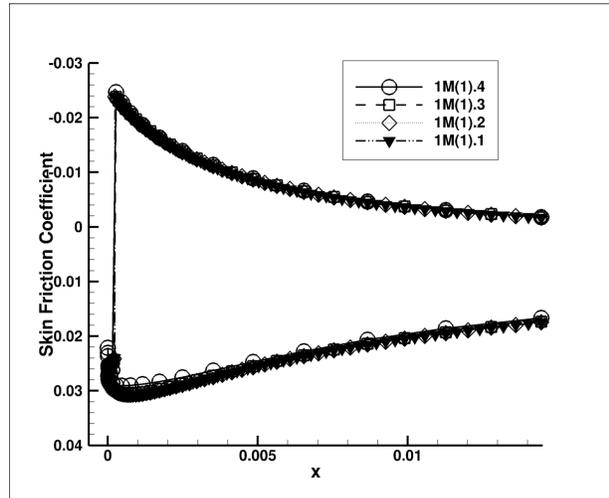


(d) Trailing edge.

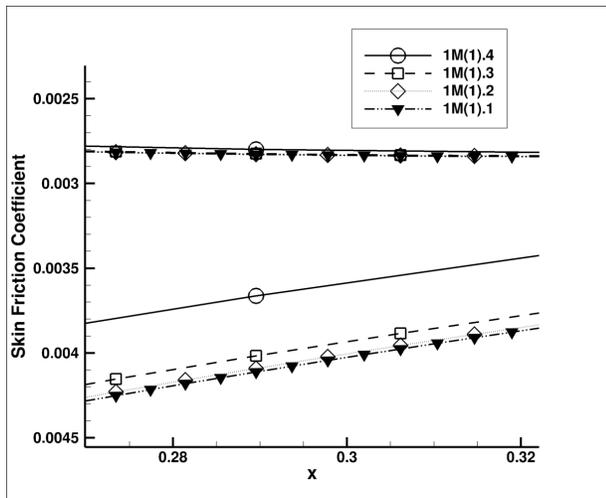
Figure A.2. The pressure coefficient, C_P , for a turbulent flow over NACA-0012 airfoil at the angle of attack of $\alpha = 10^\circ$ is computed on the 1M(1) grid family and shown at different surface locations. 1M(1) is the family of quadrilateral grids from the TMR website; the finest grid has the 1793×513 dimensions and the chord-wise mesh spacings 5×10^{-4} and 5×10^{-5} at the trailing and leading edges, respectively. The nested coarser grids in the family are derived from the finest grid by successive removal of each other grid line in each dimension. The coarse grids are identified by the number after the dot: 1M(1).1 refers to the finest grid, 1M(1).2 refers to the next coarser grid, etc.



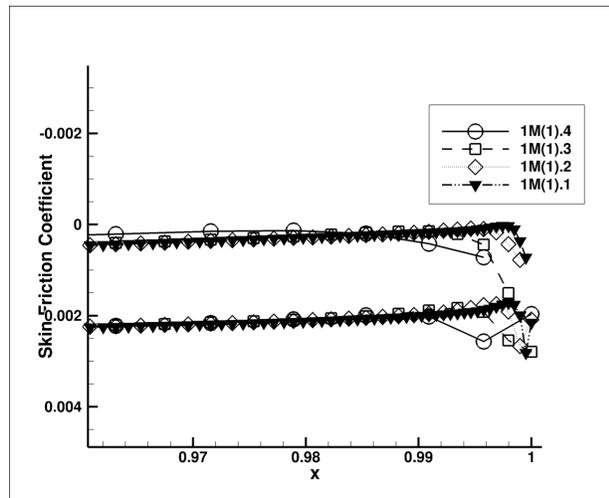
(a) Global view.



(b) Leading edge.

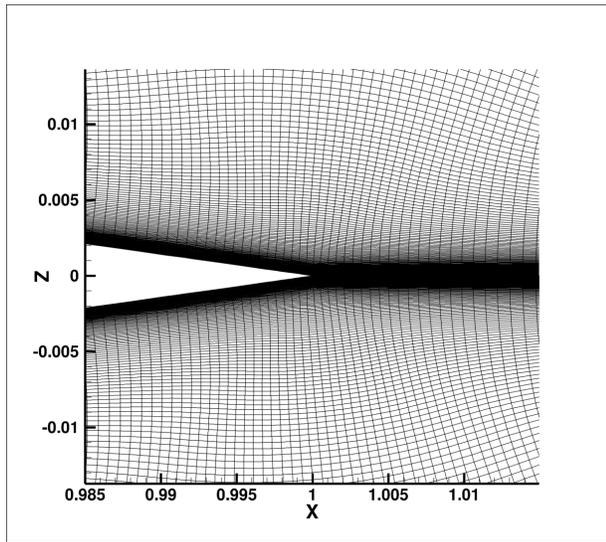


(c) Intermediate location.

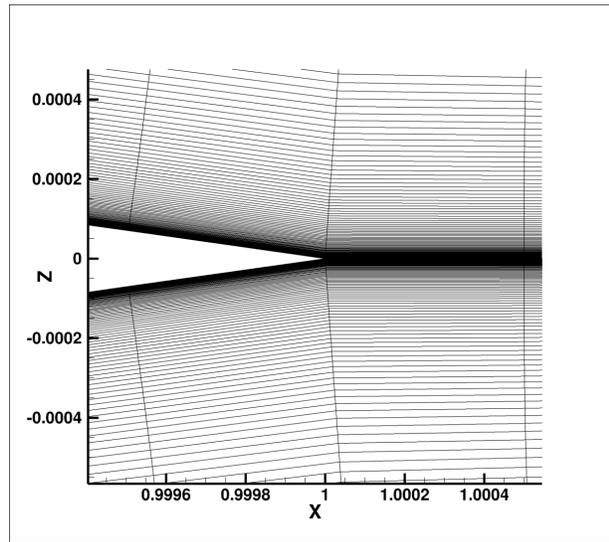


(d) Trailing edge.

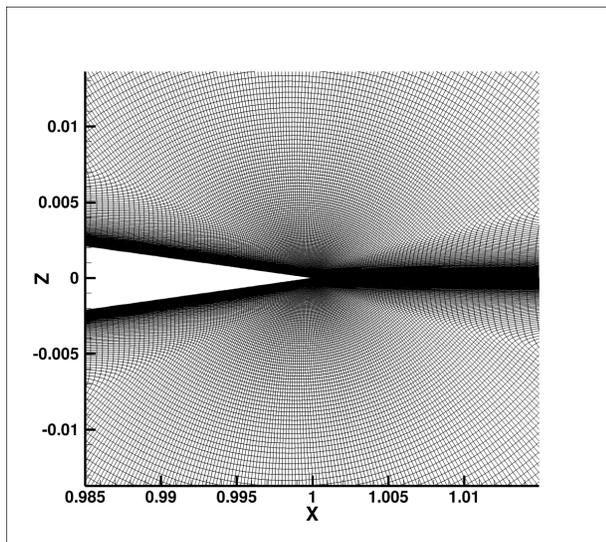
Figure A.3. The skin friction coefficient for a turbulent flow over NACA-0012 airfoil at the angle of attack of $\alpha = 10^\circ$ is computed on the 1M(1) grid family and shown at different surface locations. 1M(1) is the family of quadrilateral grids from the TMR website; the finest grid has the 1793×513 dimensions and the chord-wise mesh spacings 5×10^{-4} and 5×10^{-5} at the trailing and leading edges, respectively. The nested coarser grids in the family are derived from the finest grid by successive removal of each other grid line in each dimension. The coarse grids are identified by the number after the dot: 1M(1).1 refers to the finest grid, 1M(1).2 refers to the next coarser grid, etc.



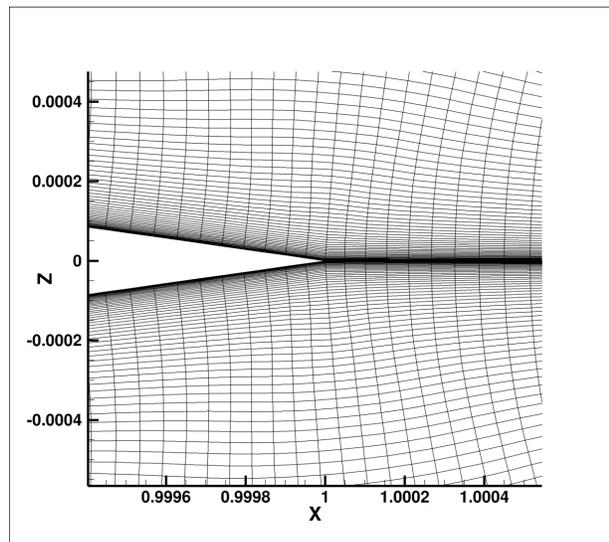
(a) 1M(1) grid, far view.



(b) 1M(1) grid, close view.

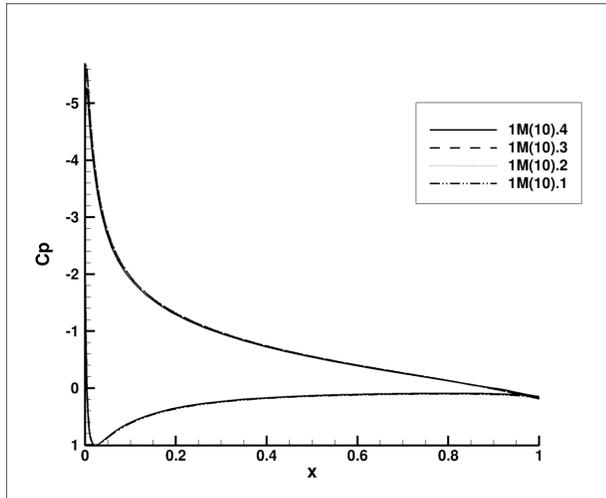


(c) 1M(10) grid, far view.

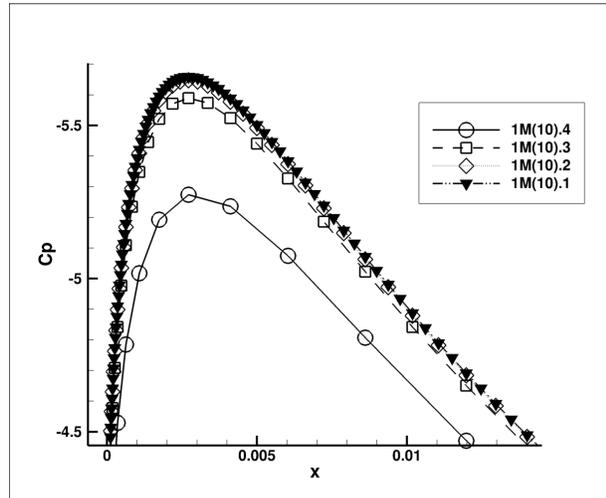


(d) 1M(10) grid, close view.

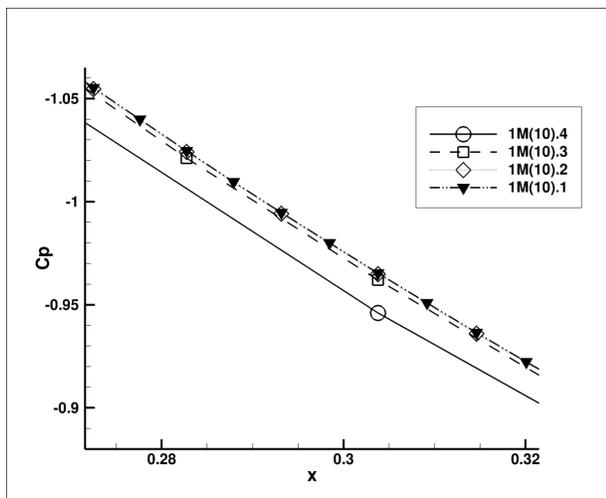
Figure A.4. Comparative views of 1M(1) and 1M(10) grids near the trailing edge on the NACA-0012 airfoil. 1M(1) is the family of grids from the TMR website; the finest grid has the 1793×513 dimensions and the chord-wise mesh spacings 5×10^{-4} and 5×10^{-5} at the trailing and leading edges, respectively. 1M(10) is a family of grids with an increased mesh density at the trailing edge; the finest grid has the same 1793×513 dimensions and the chord-wise mesh spacings 5×10^{-5} at both the trailing and leading edges.



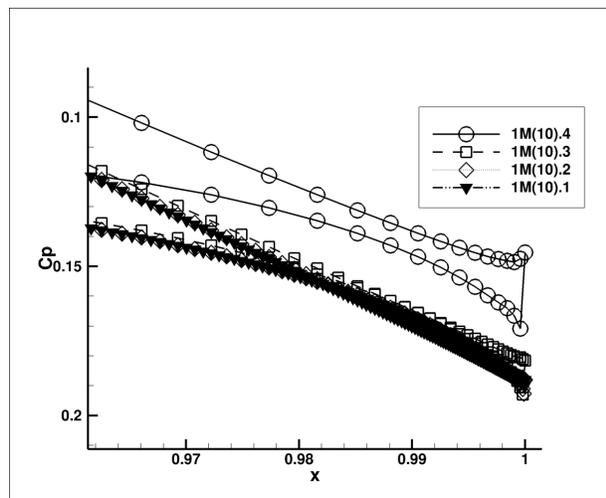
(a) Global view.



(b) Leading edge.

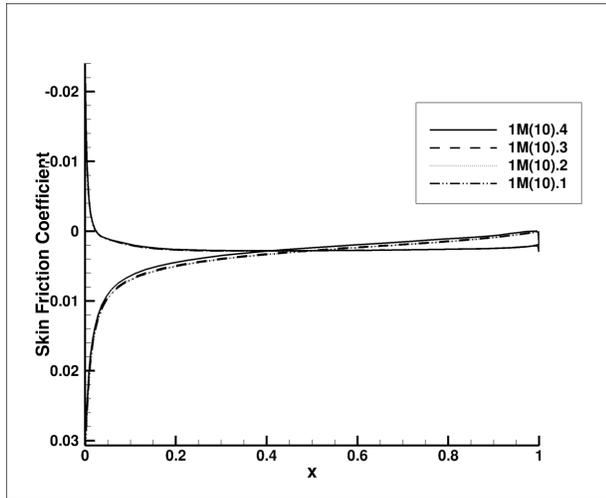


(c) Intermediate location.

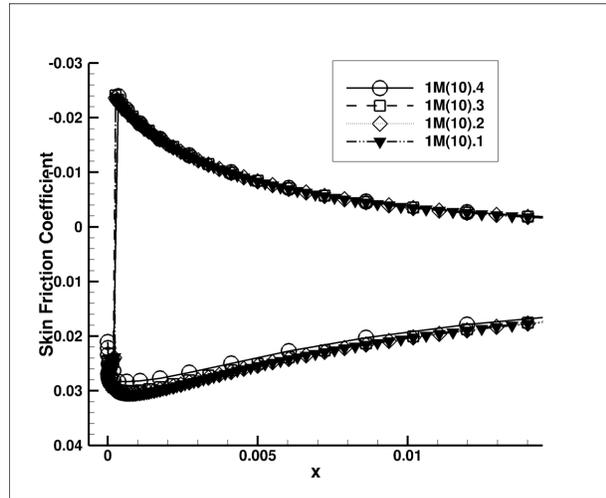


(d) Trailing edge.

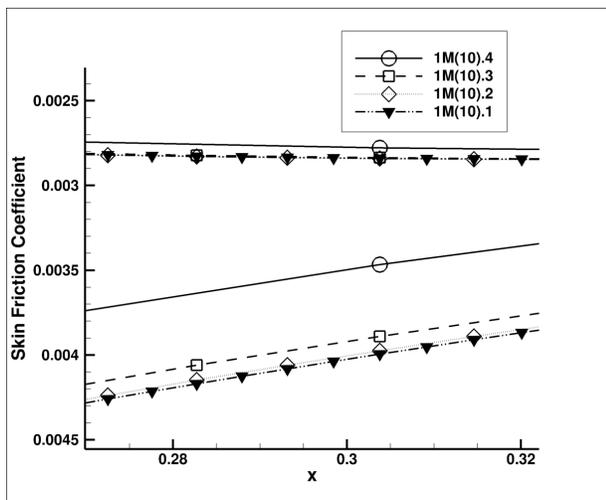
Figure A.5. The pressure coefficient, C_p , for a turbulent flow over NACA-0012 airfoil at the angle of attack of $\alpha = 10^\circ$ is computed on the 1M(10) grid family and shown at different surface locations. 1M(10) is a family of grids with an increased mesh density at the trailing edge; the finest grid has the 1793×513 dimensions and the chord-wise mesh spacings 5×10^{-5} at both the trailing and leading edges. The nested coarser grids in the family are derived from the finest grid by successive removal of each other grid line in each dimension. The coarse grids are identified by the number after the dot: 1M(10).1 refers to the finest grid, 1M(10).2 refers to the next coarser grid, etc.



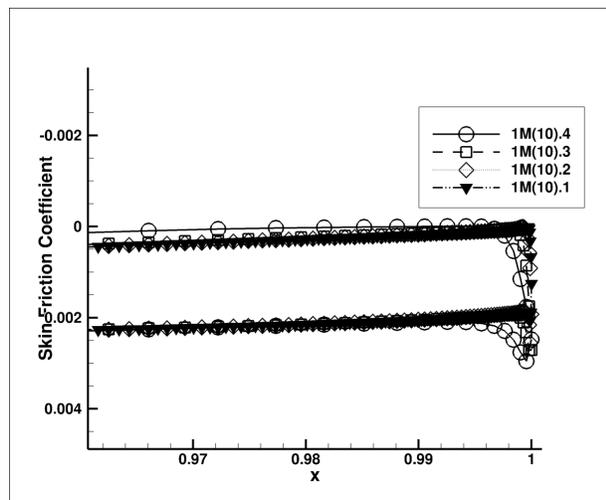
(a) Global view.



(b) Leading edge.



(c) Intermediate location.



(d) Trailing edge.

Figure A.6. The skin friction coefficient for a turbulent flow over NACA-0012 airfoil at the angle of attack of $\alpha = 10^\circ$ is computed on the 1M(10) grid family and shown at different surface locations. 1M(10) is a family of grids with an increased mesh density at the trailing edge; the finest grid has the 1793×513 dimensions and the chord-wise mesh spacings 5×10^{-5} at both the trailing and leading edges. The nested coarser grids in the family are derived from the finest grid by successive removal of each other grid line in each dimension. The coarse grids are identified by the number after the dot: 1M(10).1 refers to the finest grid, 1M(10).2 refers to the next coarser grid, etc.

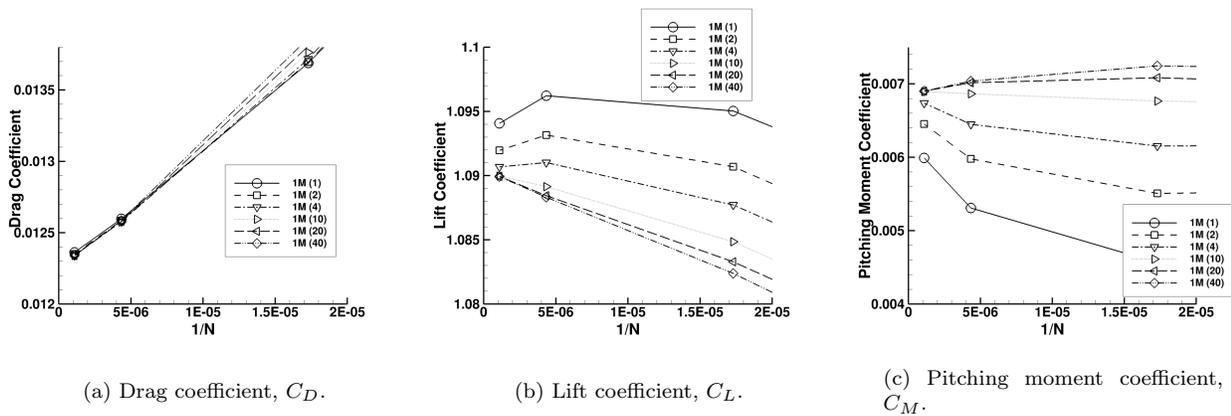


Figure A.7. The force and pitching moment coefficients for a turbulent flow over NACA-0012 airfoil at the angle of attack of $\alpha = 10^\circ$ are computed on quadrilateral 1793×513 grids with variation of the mesh density at the trailing edge. 1M(1) is the finest grid from the TMR website with the chord-wise mesh spacings 5×10^{-4} and 5×10^{-5} at the trailing and leading edges, respectively. The leading-edge mesh spacing remains the same for all grids. The number in the parenthesis indicates a reduction factor for the mesh spacings at the trailing edge relative to the 1M(1) grid; the larger number, the larger mesh density at the trailing edge.